Graph Summarization: A Survey

YIKE LIU, University of Michigan, Ann Arbor
ABHILASH DIGHE, University of Michigan, Ann Arbor
TARA SAFAVI, University of Michigan, Ann Arbor
DANAI KOUTRA, University of Michigan, Ann Arbor

arXiv:1612.04883v1 [cs.IR] 14 Dec 2016

While advances in computing resources have made processing enormous amounts of data possible, human ability to identify patterns in such data has not scaled accordingly. Thus, efficient computational methods for condensing and simplifying data are becoming vital for extracting actionable insights. In particular, while data summarization techniques have been studied extensively, only recently has summarizing interconnected data, or *graphs*, become popular. This survey is a structured, comprehensive overview of the state-of-the-art methods for summarizing graph data. We first broach the motivation behind and the challenges of graph summarization. We then categorize summarization approaches by the type of graphs taken as input and further organize each category by core methodology. Finally, we discuss applications of summarization on real-world graphs and conclude by describing some open problems in the field.

1. INTRODUCTION

As technology advances, the amount of data that we generate and our ability to collect and archive such data both increase continuously. For instance, daily activities like social media interaction, web browsing, product and service purchases, itineraries, and wellness sensors generate large amounts of data, the analysis of which can immediately impact the decision-making process and our lives. This abundance of generated data and its velocity call for data summarization, one of the main data mining tasks.

Since summarization facilitates the identification of structure and meaning in data, the data mining community has taken a strong interest in the task and, accordingly, has proposed methods for a variety of data types: sequence data and events [Casas-Garriga 2005; Kiernan and Terzi 2009; Tatti and Vreeken 2012], itemsets and association rules [Liu et al. 1999; Mampaey et al. 2011a,b; Ordonez et al. 2006; Wang and Parthasarathy 2006; Yan et al. 2005], spatial data [Lin et al. 2003], transactions and multi-modal databases [Chandola and Kumar 2005; Cordeiro et al. 2010; Shneiderman 2008; Wang and Karypis 2004; Xiang et al. 2010], data streams and time series [Cormode and Muthukrishnan 2005; Palpanas et al. 2008], video and surveillance data [Damnjanovic et al. 2008; Pan et al. 2004], and activity on social networks [Koutra et al. 2014b; Lin et al. 2008; Mehmood et al. 2013].

This survey focuses on summarizing interconnected data, otherwise known as graphs or networks. Graphs are ubiquitous, representing a broad variety of natural processes such as friendships between people, communication patterns, and interactions between chemical compounds and neurons in the brain. Formally, in its simplest form, a *plain* graph or network is an abstract data type that consists of a finite set of vertices / nodes $\mathcal{V}$ and a set of links / edges $\mathcal{E}$, the latter of which represent interactions between pairs of vertices. A graph is often represented by its adjacency matrix $\mathbf{A}$, which can be binary, corresponding to whether there exists an interaction between two vertices, or numerical, corresponding to the strength of the connection. We will refer to a graph with numerical or categorical labels/attributes for its nodes or edges as a *labeled* graph. A network that is evolving over time is called *dynamic* or *time-evolving* and is often described by a series of adjacency matrices, one per timestamp. Examples of graphs are various social networks, traffic networks, computer networks, phone call or messaging networks, location check-in networks, protein-protein interaction networks, user-product review or purchase networks, functional or structural connectomes, and more.

As the volume of interconnected data increases, summarization methods are becoming increasingly crucial. In general, *graph summarization* or *coarsening* or *aggregation* approaches seek to find a short representation of the input graph, often in the form of a summary or sparsified graph, which reveals patterns in the original data and preserves specific structural or other properties, depending on the application domain. Overall, the benefits of employing graph summarization methods can be summarized as follows:

— *Reduction of data volume and storage.* Graphs of real-world datasets are often massive; for example, the Facebook social network has over 1.7 billion users, and more than 100 billion emails are exchanged daily. Summarization techniques produce small summaries that require significantly less storage space than the original graph. Depending on the size of the summary graph, it might be possible to load it into memory, reduce the number of I/O operations, or reduce the communication volume between clusters in a distributed setting.
— *Speedup of graph algorithms and queries.* While there exist a plethora of graph methods for analyzing interconnected data, many of them cannot efficiently handle large graphs. Summarization techniques produce smaller graphs that maintain the most salient information from the original graph. The resultant summary graph can be queried and analyzed more easily and efficiently with existing tools and algorithms, and can effectively provide insights into the original data.
— *Interactive analysis support.* As the systems side makes advancements in interactive graph analysis, summarization is introduced to handle information extraction and to speed up user analysis. The resultant graph summaries make feasible the visualization of datasets that are (originally) too large to load into memory.
— *Noise elimination.* Real graph data are frequently large-scale and considerably noisy with many hidden, unobserved, or erroneous links and labels, which hinders the analysis by increasing the workload of data processing and hiding the more "important" information. Summarization is used to filter out noise and reveal the various patterns that exist in the data.

Given its advantages, graph summarization has extensive applications, including clustering [Cilibrasi and Vitányi 2005], classification [van Leeuwen et al. 2006], community detection [Chakrabarti et al. 2004], model order selection in matrix factorization [Miettinen and Vreeken 2011, 2014], outlier detection [Akoglu et al. 2012; Smets and Vreeken 2011], pattern set mining [Tatti and Vreeken 2012; Vreeken et al. 2011], finding sources of infection in large graphs [Prakash et al. 2012], understanding selected nodes in graphs [Akoglu et al. 2013], and more.

1.1. Challenges

Although we have made an attempt to give one definition of graph summarization above, the notion of a graph summary is not well-defined. A summary is application-dependent and can be defined with respect to various aspects: it can preserve specific structural patterns, focus on some entities in the network, preserve the answers to a specific set of queries, or maintain the distributions of some graph properties. In addition, graph summarization has five main challenges:

(1) *Volume of data.* The main target of graph summarization is to reduce the size of the input graph data so that other analysis can be performed efficiently. At the same time, though, summarization techniques are faced with the challenge of processing large amounts of data. The requirement of efficiency often stirs the design of summarization algorithms towards techniques that scale well with the size of the input graph.
(2) *Complexity of data.* Graph operations often cannot be easily partitioned and parallelized because of the many interactions between the entities, as well as the

complexity of entities themselves. Furthermore, the heterogeneity of nodes and edges continues to increase in real networks, and the incorporation of side information such as labels or data from heterogeneous sources (e.g. textual or image data) can require highly detailed design and quantification in algorithms. For example, in social networks, users chat or share with each other, follow or friend each other, and a single user alone has a great deal of information associated with his or her profile. Finally, real datasets often contain much noise or missing information, which may interfere with the pattern mining process.

(3) *Definition of interestingness.* Summarization involves the extraction of important or interesting information. However, the definition of "interesting" itself is subjective and identifying interesting information often requires domain knowledge as well as user preferences. Moreover, the cutoff between "interesting" and "uninteresting" can be difficult to determine in a principled way; usually it is decided by taking into account the trade-off between time, space, and information preserved in the summary, as well as the complexity of mapping solutions obtained from the summary back onto the original nodes and edges.

(4) *Evaluation.* Evaluations of summaries are manifold. From the database perspective, a summary is of good quality if it supports both global and local queries. In the context of summarizing community information, the quality of community preservation must be maximized. Compression-based techniques seek to minimize the number of bits that are needed to (losslessly) describe the input graph. Furthermore, evaluations become more complex when more elements, such as visualization and multi-resolution summaries, are involved.

(5) *Change over time.* Real data are usually dynamic, and thus graph summaries should evolve over time. For instance, social network activity, brain functions, and email communications (all naturally represented as graphs) change with time. How to incorporate the dynamic nature of such data in computation and perform analysis efficiently becomes an essential question.

As evidenced from the above-mentioned challenges, graph summarization is a challenging problem to solve in its general form. Different methods focus on different instantiations of the problem. In this survey, we categorize the methods in the literature based on the type of data they can handle, their assumptions, and the core techniques that they employ.

1.2. Our contributions

Previous work on surveying the graph summarization literature is scarce. The authors in [You et al. 2013] present some summarization algorithms for *static* graphs focusing mostly on grouping- and compression-based methods. The tutorial by Lin et al. [2013] provides more specific categorization and descriptions of ongoing work, but again only deals with static graph summarization.

Unlike these works, we cover both static *and* dynamic graph summarization, and review a wide set of proposed methodologies. We also introduce a taxonomy that can help practitioners determine the method most suitable and flexible for their data and goal (for example, compression vs. visualization). Our main contributions are:

(1) We create a taxonomy on two major types of graphs: static and dynamic graphs. Within static graphs, we separate methods into those that leverage only structural properties of graphs (i.e., they operate on plain graphs) and those that take side information into account (i.e., they handle labeled networks). Dynamic graph summarization has focused only on plain networks.

(2) We present the existing methods while highlighting properties that are useful to researchers and practitioners, such as their input/output data types and end goal.

(3) We give connections between methods of graph summarization and related fields that have potential for graph summarization, including compression, sparsification, and clustering and community detection.
(4) We review real-world applications of summarization (e.g., query efficiency, visualization), and identify open problems and opportunities for future research.

1.3. Organization

Figure 1 illustrates the taxonomy upon which we have organized this survey. Sections 2 and 3 cover summarization methods for plain and labeled *static* graphs, respectively. These methods are further categorized based on their core technique: grouping, simplification, compression, influence propagation, or pattern mining. Section 4 focuses on dynamic graph summarization, which consists of compression- and influence-based methods for plain graphs only. Section 5 presents real-world applications of graph summarization, such as query efficiency and visualization. We conclude in Section 6 by surveying some open research problems.

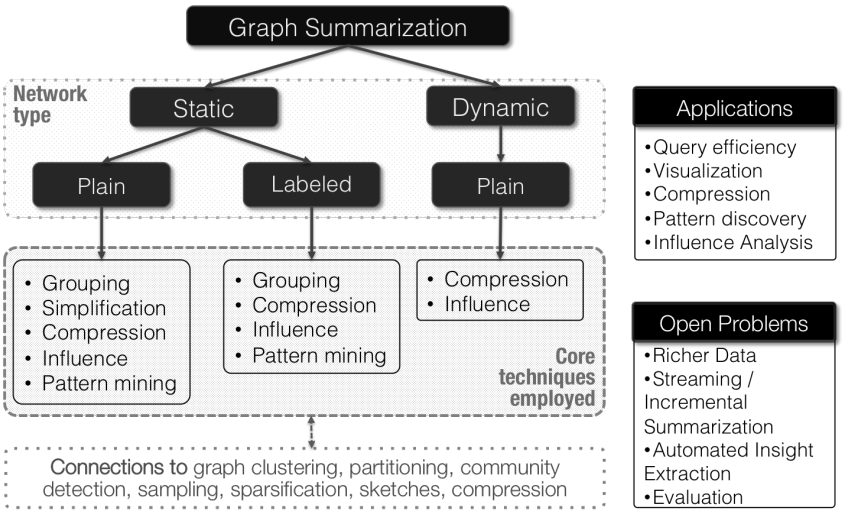

Fig. 1: *Overview of our survey: Taxonomy of graph summarization algorithms based on the input graph type and the core employed technique; approaches alternative to summarization; graph summarization applications; and open problems.*

2. STATIC GRAPH SUMMARIZATION: PLAIN NETWORKS

Most work in static graph summarization focuses solely on graph structure, i.e., plain graphs. At a high level, the problem of summarization or aggregation or coarsening of static, plain graphs is described as:

PROBLEM 1. [Summarization of Static, Plain Graphs]
Given a static graph G without any side information, or its adjacency matrix $\mathbf{A}$,
Find a summary graph or a set of structures or a compressed data structure to concisely describe the given graph.

Table I provides qualitative comparisons and explicit characterizations of static graph summarization methods for plain networks. Next, we review these methods by organizing them into categories based on the core methodology that they employ for the summarization task. When applicable, we first give the high-level idea per method type and then describe the corresponding approaches.

2.1. Grouping-based methods

Grouping-based methods are among the most popular techniques for summarization. These methods aggregate nodes into supernodes *and connect them with* superedges*, resulting in a* supergraph*. Some approaches recursively aggregate nodes/edges into supernodes/superedges based on an application-dependent optimization function, and others employ existing clustering techniques and map each densely-connected cluster to a supernode.* Given that clustering is a crucial component of these methods, we also discuss the connections between graph summarization and clustering/partitioning/community detection methods at the end of this subsection.

One representative algorithm of the grouping-based methods that leverage hierarchical clustering is GraSS [LeFevre and Terzi 2010], which targets query efficiency. This summarization method supports queries on the adjacency between two nodes, as well as the degree and the eigenvector centrality of a node. The graph summaries are generated by greedily grouping nodes so that the normalized difference of the reconstructed approximate adjacency matrix, $\bar{\mathbf{A}}$, and the original one, $\mathbf{A}$, is minimized:

$$\text{reconstructed_error} = \frac{1}{|\mathcal{V}|^2} \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{V}} |\bar{A}(i,j) - A(i,j)|$$

The resulting summaries are represented as a group of vertex sets with information about the number of edges within and between clusters. These sets are used to generate a probabilistic approximate adjacency matrix upon which incoming queries are computed: for example, if many edges cross vertex sets A and B, then it is likely that a node in A is connected to a node in B.

Focusing on graphs with weights associated to edges, the 'simple' and the 'generalized' weighted graph compression problems [Toivonen et al. 2011] merge nodes that have similar relationships to other entities (i.e., that are structurally equivalent) while minimizing approximation error and maximizing compression. The key idea is to maintain either edge weights or strengths of connections of up to a certain number of hops in the process of merging nodes to obtain the compressed graph. Specifically, simple weighted graph compression assigns each superedge the mean weight of all edges it represents, while for the generalized weighted graph compression problem, the best path between any two nodes is approximately equally good in the compressed graph as it is in the original graph, but the path does not have to be the same.

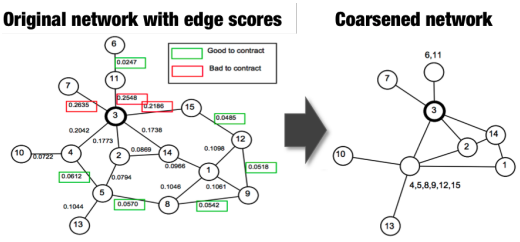

Fig. 2: COARSENET *[Purohit et al. 2014]*

An additional work that facilitates the study of diffusion and propagation processes (e.g. epidemiology, viral marketing) is COARSENET [Purohit et al. 2014], a near-linear-time algorithm that employs a scoring technique whereby a matrix perturbation argument is used to derive the change in eigenvalue due to merging two adjacent nodes such that the coarsened graph retains its diffusive properties (as characterized by the first eigenvalue of the adjacency matrix). An example of COARSENET is in Figure 2.

The previous methods group nodes based on application-specific optimization functions. However, some graph summarization algorithms simply employ existing clustering techniques. While some clustering-based summarization methods do not guarantee the quality of output, Riondato et al. [2014] proposes a method of generating supernodes and superedges with guarantees on the reconstruction error, thus improving upon methods that use heuristics, such as GraSS. Here the objective is to find the k-summary

that minimizes the l_p-reconstruction error (i.e., the p-norm of the difference between $\mathbf{A}$ and $\bar{\mathbf{A}}$) given the input graph and desired number of supernodes k. The proposed approach is the first polynomial-time approximation algorithm of its kind (with runtime $O(|\mathcal{E}| + |\mathcal{V}| \cdot k)$) to compute the "best" possible summary of a given size by employing sketching, sampling, and approximate partitioning techniques. The k-summary is guaranteed to be within a constant factor from the optimal based on solving a clustering instance (e.g., k-means) on the rows of the adjacency matrix. This method targets efficiency for the same queries as GraSS and also triangle and subgraph counting queries.

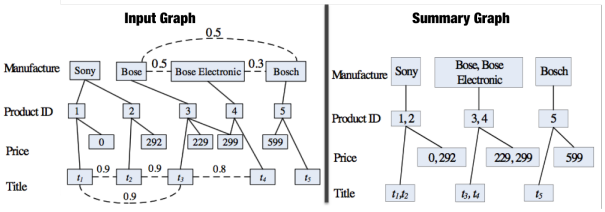

Fig. 3: *CoSum [Zhu et al. 2016]: Multi-type graph representation for a heterogeneous graph and its summary.*

Beyond the end goal of summarization itself, grouping nodes can be applied to many graph-based tasks. CoSum [Zhu et al. 2016] involves summarization on k-partite (heterogeneous) graphs for the purpose of improving entity resolution, or record linkage between data sets. Its key methodology is to transform an input k-type graph into another k-type summary graph composed of supernodes and superedges (Figure 3), using the links between instances of different types to improve the accuracy of entity resolution. CoSum jointly condenses vertices into a supernode such that each supernode is coherent (i.e., the nodes are of the same type and have high similarity), and creates superedges that connect the supernodes according to the original links between their constituent nodes. It has been demonstrated that CoSum outperforms state-of-the-art generic entity resolution approaches, especially in data sets with missing values and one-to-many or many-to-many relations.

Connection to graph partitioning, clustering, and community detection. Graph partitioning, clustering, or community detection can be used to obtain a smaller representation of an input graph by aggregating all the nodes that belong to a community into a supernode and linking them with superedges with weight equal to the sum of the cross-cluster edges, or else the sum of the weights of the original edges [Low et al. 2012; Newman and Girvan 2004; Yang and Leskovec 2013]. Figure 4 shows an example of a "summarized" clustering output. Although the clustering output can be viewed as a summary graph, a fundamental difference is that summarization methods group together nodes that are linked to the rest of the graph in a similar way, while clustering methods group together densely-connected nodes. We also note that graph clustering and relational data clustering differ significantly: graph clustering methods measure closeness based on connectivity and structural similarity, or the number of common neighbors, whereas relational data clustering is only based on attribute similarity.

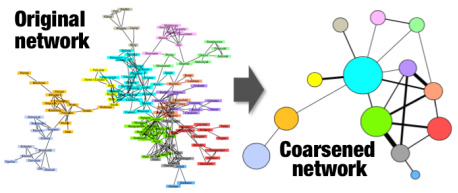

Fig. 4: *Illustration of the use of a community detection algorithm to summarize a graph (figure from [Newman and Girvan 2004]). The clustered graph with communities colored on the left is converted to the coarse-grained graph with supernodes representing the original communities.*

Comprehensive introductions to clustering techniques [Aggarwal 2015; Leskovec et al. 2014] and work on clustering or community detection methods [Aggarwal and Wang 2010] exist. Clustering methods like eigendecomposition [Shah et al. 2014], cross-associations [Chakrabarti et al. 2004], modularity-based optimization methods [Blon-

del et al. 2008; Newman and Girvan 2004; Rotta and Noack 2011], and information-theoretic co-clustering based on mutual information optimization [Dhillon et al. 2003] find tightly-knit structures like near-cliques and bipartite cores. Partitioning algorithms, such as cut-based partitioning [Dhillon et al. 2005; Karypis and Kumar 1999], spectral partitioning [Alpert et al. 1999], cross-associations [Chakrabarti et al. 2004] and their extension to hierarchies [Papadimitriou et al. 2008], and information-theoretic approaches for community detection [Rosvall and Bergstrom 2007] all seek hard clustering of all the nodes as opposed to identifying (possibly overlapping) communities.

The blockmodels representation [Faust and Wasserman 1992; Qin and Rohe 2013; Rohe et al. 2011], which comes from psychometrics and sociology and is gaining popularity in social network analysis, comprises a set of probabilistic models that combine global parameters to instantiate dense patches of connectivity (blockmodel) with local parameters that capture nodes belonging to multiple blockmodels. The stochastic blockmodel is a random-graph generative model that is often used to uncover communities in real graphs. Its original version lacks scalability, so research focuses on efficiently learning its various parameters and scaling up the process for larger networks.

2.2. Simplification-based methods

Simplification-based summarization methods streamline an input graph by removing less "important" nodes or edges, resulting in a sparsified graph. As opposed to supergraphs, here the summary graph consists of a subset of the original nodes and/or edges. In addition to simplification-based summarization methods, we also discuss existing graph algorithms that have the potential for simplification-based summarization, such as sparsification, sampling, and sketching.

A representative work on *node* simplification-based summarization techniques is OntoVis [Shen et al. 2006], a visual analytical tool that relies on node filtering for the purpose of understanding large, heterogeneous social networks in which nodes and links represent different concepts and relations, respectively. OntoVis uses information in the ontology that relates nodes and edges, such as the degree of nodes of specific type, to semantically prune the network. In addition to semantic abstraction, OntoVis allows users to do structural abstraction and importance filtering. Semantic abstraction is attained by utilizing the ontology graph such as node degree based on a specific type, and users are able to construct a derived graph from the original graph by including only nodes whose types are selected in the ontology graph. For example, in a terrorism network, selection of a node type such as terrorist organization will result in a semantic abstraction of different organizations. Structural abstraction simplifies the graph while preserving the essential structure of the entire network, for example by removing one-degree nodes and duplicate paths. Importance filtering, on the other hand, makes use of statistical measures such as node degree for evaluating connectivity and relevance between node types.

Targeting the same type of graph as OntoVis but employing edge instead of node filtering, [Li and Lin 2009] proposes an unsupervised mechanism for egocentric information abstraction, illustrated in Figure 5. The method is four-part: (i) features are automatically selected and extracted according to the surrounding network substructure, i.e., k-hop neighborhoods; (ii) the statistic dependency is measured between the features with the generation of ego nodes; (iii) distilling criteria are applied to remove less relevant information; (iv) finally, an egocentric abstracted graph is constructed incrementally, allowing the user to visualize the results. In part (iii), both frequency-based and rarity-based policies are employed to distill different kinds of information for abstraction, as "rare" and "frequent" occupy two opposite ends of the importance spectrum, both revealing potentially interesting information about a given node.

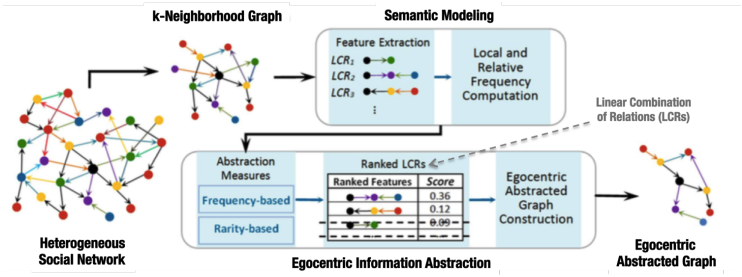

Fig. 5: *Flowchart of egocentric abstraction [Li and Lin 2009].*

Connection to graph sampling, sparsification, and sketches. An alternative method of "compressing" a graph is sampling nodes or edges from it [Hübler et al. 2008; Leskovec and Faloutsos 2006]. Graph sampling has three main goals: (i) estimation of the properties of the original graph; (ii) estimation of relative frequencies of its substructures (e.g., triangles); (iii) creation of a small subgraph that resembles the original graph so that analysis can happen on the subgraph. The third goal, which pertains most to graph summarization, is coupled with maintaining (within bounds) one or more properties of the original graph, such as the degree distribution, the size distribution of connected components, the diameter, or latent properties such as community structure [Maiya and Berger-Wolf 2010]. Various sampling techniques have been studied [Ahmed et al. 2013; Batson et al. 2013; Mathioudakis et al. 2011], and a comprehensive tutorial on graph sampling was presented at KDD [Hasan et al. 2013]. Sampling techniques include sampling edges uniformly, according to their weights or their effective resistance [Spielman and Srivastava 2008]; sampling nodes according to their in- or out-degree, PageRank, or substructures, such as spanning trees. Although sampling may allow better visualization [Rafiei and Curial 2005], it cannot detect graph structures and may need additional processing in order to make sense of the sample.

An alternative to edge sampling or sparsification is transitive reduction [Aho et al. 1972], a process that removes edges on directed graphs such that node reachability is maintained. An application of transitive reduction is described in [Dubois and Bothorel 2005], where the authors simplify the communication patterns in the Enron email network. Related in the theory community is the k-spanner, which is the sparsest subgraph in which the distance between pairs of nodes is at most k times the distance in the initial graph. A common category of the problem is the tree k-spanner, which approximates the original graph with a tree that satisfies the distance property. Finding a k-spanner is NP-hard except for the case of $k = 2$, which can be solved in $O(|\mathcal{E}| + |\mathcal{V}|)$ time. Thus, approximation algorithms have been proposed [Elkin and Peleg 2005].

Graph sketches [Ahn et al. 2012; Ghashami et al. 2016; Liberty 2013], or data synopses obtained by applying linear projections, are also relevant. Graph sketching can be viewed as linear dimensionality reduction, where the linearity of sketches makes them applicable to the analysis of streaming graphs with node and edge additions and deletions and distributed settings, such as MapReduce [Dean and Ghemawat 2004].

2.3. Compression-based methods

Compression is a common technique in data summarization. Many approaches cast the summarization problem as a compression problem where the goal is to minimize the number of bits needed to describe the input graph via its summary, which can be seen as a model for the input graph, and its unmodeled parts. The graph summary or model is significantly smaller than the original graph, and often reveals various structural patterns. However, there are also methods that aim at solely compressing a given graph without necessarily creating a graph summary. Here we focus mostly on the former methods, which often view the summarization problem as a model selection task and

employ the two-part Minimum Description Length (MDL) code, the goal of which is to minimize the description (number of bits) of the given graph G and the model class M:

$$minL(G, M) = L(M) + L(G|M), \tag{1}$$

which is given as the description length of the model, $L(M)$, and the description length of the graph given the model (i.e., the errors or unmodeled parts with respect to the model). For completeness, we also present some graph compression methods which could be adapted to summarization, although not originally designed for that purpose.

Relying on MDL, [Navlakha et al. 2008a] proposes a highly compact two-part representation, using the summaries to compress graphs with bounded error. This representation, which is obtained by aggregating nodes in the summary generation, consists of a graph summary S and a set of corrections C (example given in Figure 6). The summary is an aggregate graph in which each node corresponds to a set of nodes in G, and each edge represents the edges between all pairs of nodes in the two sets. The correction term specifies the list of edge-corrections that should be applied to the summary to exactly recreate G. The cost of a representation, R, is the sum of the storage costs of both S and C: $cost(R) = |E_S| + |C|$, where E_S is the set of superedges in S. To find the MDL representation, two algorithms are presented: GREEDY and RANDOMIZED, both of which aggregate groups of nodes (thereby falling into the grouping-based summarization category) and check if the MDL cost of the graph decreases. GREEDY iteratively combines node pairs that give the maximum cost reduction into supernodes, whereas RANDOMIZED randomly picks a node and merges it with the best node in its vicinity.

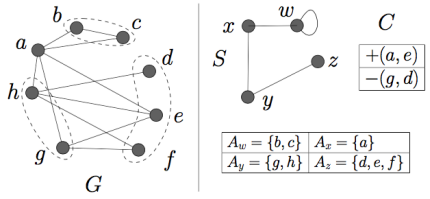

Fig. 6: *Example of MDL representation [Navlakha et al. 2008a]: S and C are its two parts. The bottom-right table shows the supernode mapping.*

Addressing an information-theoretic optimization problem also based on MDL, VOG [Koutra et al. 2014b], or *vocabulary-based summarization of graphs*, aims to succinctly describe large-scale graphs with a few, possibly overlapping and easily understood structures, which are encoded in the model M. The graph summary is given in terms of a predefined "vocabulary" of structures that goes beyond the simple rectangles that most summarization and clustering methods find, identifying cliques and near-cliques, stars, chains, and (near-) bipartite cores. VOG is modular (Figure 7): (i) it first performs graph clustering by adapting the node reordering method SLASHBURN [Lim et al. 2014] to extract egonetworks and other disconnected components; (ii) then labels the extracted subgraphs with the appropriate structures in the assumed vocabulary using MDL as a model selection criterion; and (iii) finally creates a summary by employing heuristics that choose only the subgraphs that minimize the total encoding cost of the graph, $L(G, M)$, as it is defined in Equation (1).

A VOG variant called VOG-OVERLAP has been proposed to compare the summarization power of state-of-the-art clustering methods (including METIS, LOUVAIN, SPECTRAL CLUSTERING, and SLASHBURN), which are ultimately used in the subgraph generation step of VOG [Liu et al. 2015]. VOG-OVERLAP summarizes graphs with important structures that lead to small edge overlap and large node/edge coverage. Furthermore, similar to VOG, [Miettinen and Vreeken 2011, 2014] discuss MDL for Boolean matrix factorization, which can be viewed as a summary in terms of possibly overlapping full cliques in directed graphs.

Beyond MDL, Graph Dedensification [Maccioni and Abadi 2016] is a technique that losslessly compresses the neighborhood around high-degree nodes, accelerating query

processing and enabling direct operations over the compressed graph. Based on the intuition that a large amount of information redundancy surrounds high-degree nodes and thus can be synthesized and eliminated, Graph Dedensification introduces compressor nodes that represent common connections of a set of nodes to high-degree ones, so that many redundant edges can be removed from the original graph (Figure 8). Specifically, the method constrains the formation of compressor nodes in order to provide global guarantees about the structure of the compressed graph. In other words, constraints are placed on how and when dedensification occurs, so that the query executer is able to leverage them in order to reduce the scope of compressor handling during query processing. The dedensification only creates a new compressor node if a constraint holds on a set of high-degree nodes. The guarantees are then used to create query processing algorithms that enable direct querying, in particular pattern matching queries, of the compressed graph. Dedensification is similar to virtual node compression from the Web search community [Buehrer and Chellapilla 2008] (Section 2.5), which mainly targets community-based queries, as well as some ideas in visualization-based summarization [Dunne and Shneiderman 2013]. Unlike approximate querying approaches, Graph Dedensification provides exact answers due to its losslessness. Furthermore, it does not suffer from the space/time trade-off of graph indexing.

A representative work on compressed graph querying provides a blueprint for performing lossless queries after compressing graphs, achieved by defining functions for graph compression, query conversion, and result interpretation [Fan et al. 2012]. For example, this blueprint can be implemented for queries of reachability (can node A be reached from node B?) and pattern matching (is there a subgraph that best satisfies a function provided by the user on path length between nodes in the subgraph?). This work is also adapted to handle incremental changes in the original graph by reflecting them in the compressed representation without needing to decompress it.

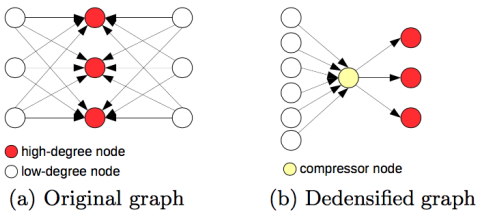

Fig. 8: *Example of Graph Dedensification [Maccioni and Abadi 2016]: edges are significantly removed after the addition of compressor node.*

Connection to graph compression. Graph summarization and compression are closely related. Graph compression can be used to obtain summaries (i.e., compressed graph representations in this case), which reduce the storage space and/or increase query efficiency.

To effectively compress social networks for neighbor queries, or queries searching for all neighbors of a vertex, [Maserrat and Pei 2010] proposes a lossless compression-based approach employing a novel Eulerian data structure using multi-position linearizations of directed graphs. This method is the first to answer both out-neighbor and in-neighbor queries in sublin-

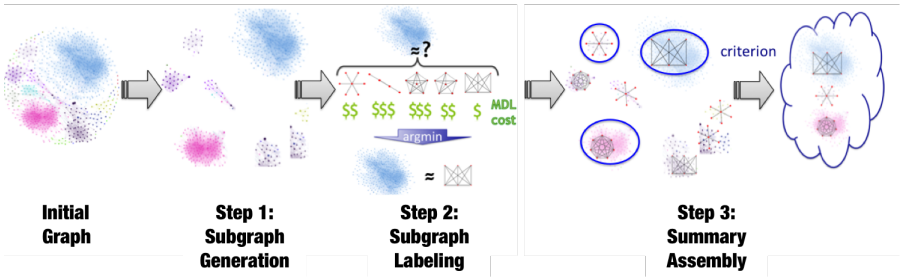

Fig. 7: VoG *[Koutra et al. 2014b]: overview of vocabulary-based graph summarization.*

ear time.This work is extended to the problem of community-preserving lossy compression of social networks [Maserrat and Pei 2012], where a sequence graph compression approach is proposed with a greedy algorithm that optimizes the objective function.

On the systems side, GBASE [2011], a billion-scale graph management system, is capable of managing and mining graphs with a lossless compression scheme presented as zip block encoding. GBASE, which is designed for parallelism, supports efficient global and targeted queries. In the indexing stage, the original raw graph is partitioned into blocks with which nodes are reshuffled to put those in the same partition nearby. Next, non-empty blocks are compressed and stored with their meta-information. In the query stage, GBASE supports efficient global and targeted queries.

The bipartite minimum logarithmic arrangement (BIMLOGA) problem [Dhulipala et al. 2016] extends the techniques of graph reordering and bisection to improve the compression of graphs and inverted indices, finding the best "compression-friendly" order for an input graph or an inverted index. BIMLOGA finds an optimal permutation of nodes using an objective function that is closely related to minimizing the number of bits needed to store a graph or index represented with the delta-encoding scheme.

Other works on graph compression include: compression of web graphs using lexicographic localities [Boldi and Vigna 2004b]; extensions to social networks [Chierichetti et al. 2009; Grabowski and Bieniecki 2014]; breadth first search-based approaches [Apostolico and Drovandi 2009]; lossy edge encoding per triangle [Feng et al. 2013]; node reordering [Lim et al. 2014] in real graphs with power-law degree distributions; weighted graph compression so that the edge weights of up to a certain number of hops are maintained [Toivonen et al. 2011]; and provably optimal compression of Erdös-Rényi random graphs using structural entropy (SZIP [Choi and Szpankowski 2012]).

There is also extensive summarization work on general data compression, some of which has the potential of application on graphs. For example, in order to effectively summarize uncertain transaction data, the Minimal Probabilistic Tile Cover Mining [Liu and Chen 2016] method aims to find a high-quality probabilistic tile set covering an uncertain database with minimal cost. This is essentially a compression problem, as it generates a summary of transaction data with a more concise and interpretative representation by optimizing the compression cost. Applications of such transaction data exist in other data representations, such as binary matrices and bipartite graphs.

2.4. Influence-based methods

Influence-based summarization methods aim to discover a short representation of the influence flow in large-scale graphs. Usually such methods formulate the summarization problem as an optimization process in which some quantity related to information influence is maintained.

Social graph summarization via information propagation and social influence analysis is the focus of Community-level Social Influence, or CSI [Mehmood et al. 2013]. CSI detects a set of communities using METIS and their reciprocal influence by extending the Independent Cascade model [Kempe et al. 2003] to communities instead of individual nodes. To balance between data fit and model complexity, the method uses the MDL and Bayesian Information Criterion (BIC) approaches for model selection. Unlike influence propagation approaches that find representative cascades for information diffusion, CSI leads to a compact representation of the input network where the nodes correspond to communities and the directed edges represent influence relationships. This is also different from grouping-based summarization techniques where the superedges simply represent the aggregate connections between the adjacent supernodes.

SPINE [Mathioudakis et al. 2011], an alternative to CSI, is based on the sparsification of influence networks. Its goal is to eliminate a number of edges in the input network in order to keep only the ones that "explain" the information propagation, i.e.,

those that maximize the likelihood of the observed data. Because the problem is NP-hard, SPINE is greedy and thus leads to a slight quality decrease. Unlike CSI, SPINE simply eliminates edges and does not group nodes into communities or supernodes.

2.5. Pattern-mining-based summarization

Pattern mining techniques aim to summarize an input network via structural patterns. Some methods in this category replace frequent patterns, such as bipartite cores, with supernodes or motifs, thus reducing the size of the original graph and enabling its visualization. Other methods represent a network with its anomalous patterns or 'summarized' distribution plots of graph properties.

Virtual Node Mining (VNM) [Buehrer and Chellapilla 2008] is a lossy compression scheme for the Web graph, accommodating community-related queries and other random access algorithms on link servers. Like SUBDUE [Cook and Holder 1994] (discussed in Section 3.2), VNM uses a frequent mining approach to extract meaningful connectivity formations by casting the outlinks/inlinks of each vertex as a transaction/itemset. Then, for each recurring pattern, it removes the links from its vertices and generates a new vertex (virtual node) in the graph, which is added as an outlink (Figure 9). VNM facilitates global mining, enabling discovery for both inter-domain (e.g. two competing companies do not link to each other, but third parties will often link to both) and intra-domain (e.g. page belongs to a community) patterns, hence allowing incremental graph updates. The complexity of VNM's mining phase is $O(|\mathcal{E}| \log |\mathcal{E}|)$.

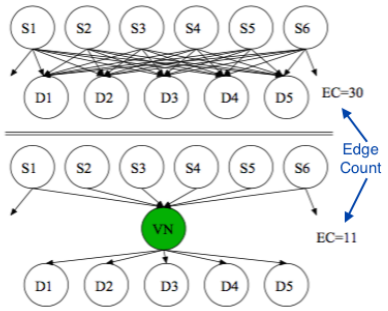

Fig. 9: *VNM [Buehrer and Chellapilla 2008] on a bipartite graph.*

Relatedly, the idea of motif simplification [Dunne and Shneiderman 2013] is introduced to enhance network visualization, where common patterns of nodes and links are replaced with compact glyphs to help visualize and understand the complex relationships between entities and their attributes. The idea of a virtual or compressor node in VNM and dedensification [Maccioni and Abadi 2016], respectively, are related to the 'connector' motif in this work. Each method introduces a different algorithm for detecting these patterns, but the nodes that they introduce are equivalent. Some of these motifs are part of the vocabulary used in VoG [Koutra et al. 2014b], which also allows for the near-structures that appear often in real-world graphs.

Other works in pattern-mining based summarization focus on rescaling a visualized graph or carefully selecting the graph components that are needed for visualization. For example, Apolo [Chau et al. 2011] helps with attention routing in graphs by summarizing the graph with the neighborhoods of a few user-selected seed nodes. Additionally, these neighborhoods can be interactively explored. A follow-up anomaly detection system, OPAvion [Akoglu* et al. 2012], mines graph features using the Hadoop-based graph mining framework Pegasus [Kang et al. 2009], spots anomalies by employing OddBall [Akoglu et al. 2010] for mining distributions of egonet-related features, and interactively visualizes the anomalous nodes via Apolo. These two systems focus on summarizing the graph based on user selections or anomalous nodes that are detected with dedicated algorithms. PERSEUS [Koutra et al. 2015] takes a different approach to summarization; it is a large-scale system that enables the comprehensive analysis of large graphs by supporting the coupled summarization of graph properties (computed via Pegasus) and structures, guiding attention to outliers, and allowing the user to interactively explore normal and anomalous node behaviors in distribution plots and egonetwork representations. Figure 10 gives an overview of the front-end of PERSEUS.

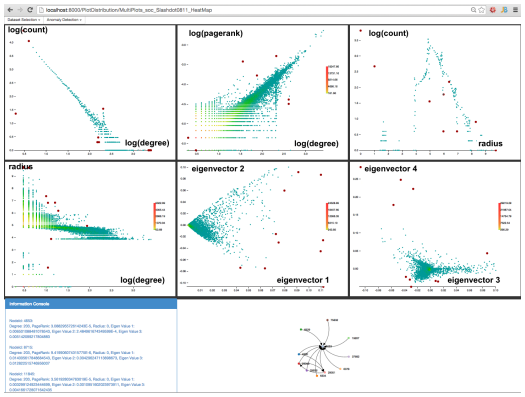

Fig. 10: *The front-end of* PERSEUS*, with linked plots for graph properties and a dynamic egonet. The annotated red points correspond to density-based anomalies discovered during offline pre-processing.*

Other "summarization" methods include scaled density plots to visualize scatter plots [Shneiderman 2008], random and density sampling [Bertini and Santucci 2004] and rescaled visualization of spy, distribution and correlation plots of massive graphs [Kang et al. 2014].

Table I: *Qualitative comparison of static graph summarization techniques. The first six columns describe the type of the input graph (e.g. with weighted / directed edges, and one / multiple types of node entities), followed by three algorithm-specific properties (i.e., user parameters, algorithmic compexity—linear on the number of edges or higher—, and type of output). The last column gives the final purpose of each approach. Notation: (1) * indicates that the algorithm can be extended to handle the corresponding type of input, but the authors do not provide details in the paper, for complexity * indicates sub-linear; (2) + means that at least one parameter can be set by the user, but it is not required (i.e., the algorithm provides a default value).*

	Input Graph						Algorithmic Properties				
Method	**Unweighted**	**Weighted**	**Undirected**	**Directed**	**Homogeneous**	**Heterogeneous**	**Parameter-free**	**Linear**	**Core Technique**	**Output**	**Objective**
GraSS [LeFevre and Terzi 2010]	✔	✘	✔	✘	✔	✘	✘	✘	grouping	supergraph	query efficiency
Weighted Graph Compression [Toivonen et al. 2011]	✘	✔	✔	✔*	✔	✘	✘	✘	grouping	supergraph	compression
COARSENET[Purohit et al. 2014]	✘	✔	✘	✔	✔	✘	✘	✔*	grouping	supergraph	influence
l_p-reconstr. Error [Riondato et al. 2014]	✔	✔	✔	✘	✔	✘	✘	✔	grouping	supergraph	query efficiency
CoSum [Zhu et al. 2016]	✔	✔*[1]	✔	✘	✘	✔	✔	✔	grouping	supergraph	entity resolution
OntoVis [Shen et al. 2006]	✔	✘	✔	✘	✘	✔	✔	✔	simplification	sparsified graph	visualization
Egocentric Abstraction [Li and Lin 2009]	✔	✘	✘	✔	✘	✔	✘	✘	simplification	sparsified graph	influence
MDL Representation [Navlakha et al. 2008a]	✔	✘	✔	✔*	✔	✘	✔	✘	compression	supergraph	compression
VOG[Koutra et al. 2014b]	✔	✘	✔	✘	✔	✘	✔	✔*	compression	structure list	patterns, visualization
Graph Dedensification [Maccioni and Abadi 2016]	✔	✘	✔*	✔	✘	✔	✔	✔*	compression	sparsified graph	query efficiency
Query-Preserving Compression [Fan et al. 2012]	✔	✘	✘	✔	✔	✘	✘	✔	compression	supergraph	query efficiency
Eulerian Data Structure [Maserrat and Pei 2010]	✔	✘	✘	✔	✔	✘	✘	✔	compression	graph linearization	compression
Community-Preserving Compr. [Maserrat and Pei 2012]	✔	✘	✔	✔*	✔	✘	✘	✘	compression	graph linearization	compression
GBASE[Kang et al. 2011]	✔	✔	✔	✔	✔	✘	✘	✔*[2]	compression	compressed graph	query eff., compression
Graph Bisection [Dhulipala et al. 2016]	✔	✘	✔	✘	✔	✘	✔	✔*	compression	compressed graph	compression
SZIP [Choi and Szpankowski 2012]	✔	✔	✔	✘	✔	✘	✘	✔	compression	compressed graph	compression
BV and BL Schemes [Chierichetti et al. 2009]	✔	✘	✘	✔	✔	✘	✘	✘	compression	compressed graph	compression
CSI [Mehmood et al. 2013]	✔	✘	✘	✔	✔	✘	✔	✘	influence	supergraph	influence
SPINE [Mathioudakis et al. 2011]	✘	✔	✘	✔	✔	✘	✘	✘	influence	sparsified graph	influence
VNM [Buehrer and Chellapilla 2008]	✔	✘	✘	✔	✔	✘	✘	✘	influence	sparsified graph	compression
Motif Simplification [Dunne and Shneiderman 2013]	✔	✘	✔	✘	✔	✔	✘	✘	pattern-mining	supergraph	visualization

3. STATIC GRAPH SUMMARIZATION: LABELED NETWORKS

So far we have only studied summarization methods based on the structural properties of static graphs, ignoring additional information like node and edge attributes. However, in reality, graphs are often annotated, labeled, or attributed. For example, in a social network, a typical node representing a user will contain information about age, gender, and location; transportation graphs may have information about the capacity of edges and the maximum speed per street; forums like Quora, which can be interpreted as networks of questions and answers, have comments, upvotes, and downvotes.

Currently, all existing works focus on node attributes alone, so we overlook other types of side information, although they are certainly of interest in summarization tasks. For instance, joint summarization of heterogeneous sources of data (including graphs, textual data, streaming data sources) may have various applications, but are more challenging and are underexplored in the literature. A formal definition of graph summarization for static, labeled graphs is given below:

PROBLEM 2.[Summarization of Static, Labeled Graphs]
Given a static graph G, or, equivalently, its corresponding adjacency matrix $\mathbf{A}$, and side information, such as node attributes
Find a summary graph or a set of labeled structures or a compressed data structure to **concisely describe** the given graph.

Apart from inheriting the challenges of summarizing static graph structure, the main challenge in summarizing labeled graphs is the efficient combination of two different types of data: structural connections and attributes.

Table II provides qualitative comparisons and explicit characterizations of static graph summarization methods for labeled graphs, which we review next by classifying them based on their core technical methodology. The overview of this section is included in Figure 1, which provides the taxonomy that we introduce in this survey.

3.1. Grouping-based methods

Grouping-based methods aggregate nodes into supernodes connected by superedges based on both structural properties and node attributes. This is done by grouping nodes that are structurally close (i.e., within a few hops) and share similar attribute values.

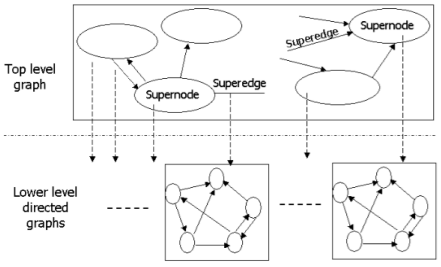

Fig. 11: *S-Node representation [Raghavan and Garcia-Molina 2003] of a Web graph.*

Optimizing specifically for Web graphs, S-Node representation [Raghavan and Garcia-Molina 2003] is a novel two-level scheme that achieves the lossless compression of such graphs. Here a Web graph is a set of small directed graphs consisting of supernodes and superedges, which are pointers to lower-level graphs that encode the interconnections within a small subset of Web pages (Figure 11). By exploiting empirically observed properties of Web graphs (e.g. link copying, domain and URL locality, page similarity) to guide the grouping of pages into supernodes and using a compression technique called reference encoding for the lower level directed graphs, S-Node representation achieves high space efficiency and naturally isolates portions of Web graphs relevant to particular queries. S-Node representa-

[1]Weights only exist between nodes of the same type.
[2]Linearity of GBASE depends on the clustering method used.

tion is the first Web graph representation scheme to combine compression with support for both complex queries and local graph navigation.

SNAP and k-SNAP are two popular database-style operations for summarizing graphs [Tian et al. 2008]. The SNAP algorithm, which serves as a foundation for the k-SNAP evaluation heuristics, begins by grouping nodes based only on the attributes, and then iteratively splits groups in the current grouping until the grouping is compatible with the relationships, eventually producing the maximum attribute- and relationship-compatible grouping (Figure 12). The nodes of the summary graph given by SNAP correspond to the groups, and the edges are the group relationships inferred from the node relationships within the selected edge types. k-SNAP further allows users to control the summary resolution and provides the "drill-down" and "roll-up" abilities to navigate through summaries with different resolutions. The authors also propose an efficient algorithm to evaluate the SNAP operation, as well as heuristics to approximate the k-SNAP operation, which is proved to be NP-complete.

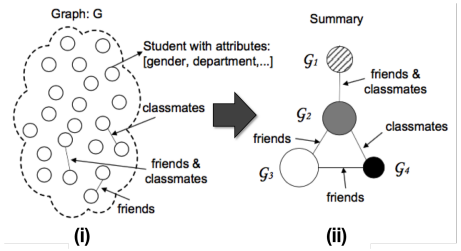

Fig. 12: *Graph summarization by aggregation of SNAP [Tian et al. 2008]: (i) graph about students; and (ii) its summary graph.*

An automation of this method where the user input aspect is removed is given by CANAL [Zhang et al. 2010]. In order to make interactive summarization more useful, CANAL automatically categorizes numerical attribute values by exploiting domain knowledge about the node labels and graph structure. To point users to the potentially most insightful summaries, the method incorporates three tenets of "interestingness": (i) *Diversity*, the number of strong relationships connecting groups with different attribute values; (ii) *Coverage*, the fraction of nodes in the original graph that are present in strong group relationships; (iii) *Conciseness*, the sum of the number of groups and strong group relationships, where a lower sum is preferred. Overall, interestingness is given as $\frac{Diversity(S) \times Coverage(S)}{Conciseness(S)}$, where S is the summary graph.

Another database-centered approach to graph summarization is given in [Hassanlou et al. 2013], where the summary is based on grouping together nodes that have the *same* attribute values using the GROUP BY operation. The proposed method is also applied to probabilistic graphs, which are graphs with edges that have probabilities of existence associated with them. The extended method is based on calculating the expected value of the summaries rather than their exact values. Shoaran et al. [2013] extends this work by aiming to protect the privacy of data in the labeled summaries generated by the probabilistic technique. Proposed is a privacy framework that extends upon Zero-Knowledge Privacy [Gehrke et al. 2003], which improves on differential privacy by only considering a random sampling of data with added noise for the summarization. This technique is also expanded to work with probabilistic attributed graphs by adding noise to the expected quantities of the sampled graph before summarization.

Connection to graph clustering, and community detection. There are many clustering and community detection methods for labeled networks, a few of which we review here. These do not perform summarization, but could be leveraged by summarization approaches in order to obtain compact representations of graphs with attributes. As we mention in the case of plain graphs, one fundamental difference between summarization and clustering is that the former finds coherent sets of nodes with similar connectivity patterns to the rest of the graph, while clustering aims at discovering coherent, densely-connected groups of nodes. For instance, in [Zhou et al. 2009], the goal is to find cohesive groups of nodes with homogeneous vertex properties by balancing structural

and attribute similarities. An augmented graph is generated by adding extra nodes for each value an attribute can take and connecting the original nodes to the new nodes with the corresponding attribute values. Random Walk with Restarts [Tong et al. 2006] is used on this augmented graph to form the distance matrix, which is plugged into the k-medoids algorithm. Also provided is a method to give each attribute a different weight, which can be optimized for the best clustering. Yang et al. [2013] tackles the attributed graph clustering problem by expanding on the authors' previous method [Yang and Leskovec 2013], which works only on plain graphs. They create a generative model by assuming a distribution for the cluster memberships which then generates the edges and attributes. The clusters are then assigned using maximum likelihood estimation. A similar generative model is used in [Xu et al. 2012], which assigns the clusters using a variational algorithm (a method for solving intractable Bayesian inference problems).

3.2. Compression-based methods

Most compression-based summarization methods leverage the MDL principle to guide the grouping of nodes or the discovery of frequent structures to be replaced with virtual nodes in the summary. In addition to the structure, which is used in static plain graph summarization, in this case, the compression and/or aggregation techniques leverage both the node structure and attributes.

SUBDUE [Cook and Holder 1994] is one of the first and most famous frequent-subgraph-based summarization schemes. It employs a two-part MDL representation, which we describe in Section 2.3, and greedy beam search to iteratively replace the most frequent subgraph in a labeled graph (i.e., the substructure that minimizes the MDL cost) with a meta-node, and, thus, discovers small lossy descriptions of labeled graphs. Multiple passes of SUBDUE eventually produce a hierarchical description of the structural regularities in the graph. The resulting representation can be used to either identify anomalous structures (i.e., instances that do not compress well) or the most common substructures (i.e., substructures that have very low compression cost). Since the introduction of SUBDUE, methods have been proposed to alleviate the complexity issues of frequent pattern mining on graphs, or else extend its application in different settings: [Maruhashi et al. 2011] proposes *MultiAspectForensics*, a tool to detect and visualize graph patterns; [Thomas et al. 2010] introduces MARGIN, an algorithm that reduces the search space of frequent subgraphs by only mining the maximal frequent subgraphs of a graph database; and [Wackersreuther et al. 2010] proposes a frequent subgraph mining algorithm to operate on dynamic graphs.

Wu et al. [2014] also models attributed graph summarization as an information-theoretic problem, using MDL to minimize the sum of the model description size and the size of the data with the help of the model. The model cost consists of the number of bits to describe three parts: the number of node and attribute groups; the nodes in each group; and the links among groups. The data cost consists the description cost of the links inside each group and the attributes. A greedy algorithm generates summaries of labeled graphs where the MDL cost is used as a measure of whether a certain node grouping is beneficial to the summary as a whole (i.e., it reduces the total encoding cost of the graph). A faster version of the greedy algorithm is obtained by initializing the summaries using label propagation instead of random initialization.

Khan et al. [2014] leverages Locality Sensitive Hashing (LSH) and MDL to compress large networks for the purposes of in-memory processing and summary generation. LSH is a popular method for avoiding the quadratic number of pairwise comparisons when measuring the similarity between entities (e.g. nodes in the graph setting). In this work, LSH works on graphs by performing rounds of computing minhash functions on node neighborhoods, combining these minhash functions into groups, computing hash codes on the groups, and then aggregating the nodes that have the same hash

codes (i.e., the nodes that fall into the same "buckets" based on the outputs of multiple hash functions). To handle the labeled graphs, the proposed method concatenates the adjacency and attribute lists together before performing LSH. Supernodes are used to combine nodes—unlike other works, virtual nodes are used to combine edges between groups of nodes. MDL is used to measure the relative increases in compression efficiency achieved by grouping nodes to supernodes and edges to superedges, and, thus guiding the summarization process.

We further note that MDL is used frequently for data that, while not explicitly modeled as a graph, can be implicitly viewed as such: R-KRIMP and RDB-KRIMP [Koopman and Siebes 2008, 2009] summarize multi-relational data, which can be viewed as attributed graphs. The former algorithm, R-KRIMP, finds characteristic patterns in single data tables and then with this reduced search space finds a small set of multi-relational characteristic item sets. The latter extends upon R-KRIMP by finding more expressive and thus more descriptive patterns.

3.3. Influence-based methods

Currently influence-based summarization methods for labeled graphs are scarce. The representative method in this category leverages both structural and node attribute similarities in the problem formulation in order to summarize the influence or diffusion process in a network.

VEGAS [Shi, Tong, Tang, and Lin Shi et al.] is a recently-proposed matrix decomposition-based algorithm for understanding influence in citation networks. The summarization problem is formulated as finding the community membership matrix H of the nodes (i.e., papers in the citation network) such that $\min_{H \geq 0} ||M^G - HH^T||_F^2$, where $M^G = \frac{AA^T + A^TA}{2}$ is the node similarity matrix and A is the adjacency matrix. In the case of labeled networks, M^G is replaced with the generalized similarity matrix $M^D = \frac{(A \odot A^D)(A \odot A^D)^T + (A \odot A^D)^T(A \odot A^D)}{2}$ to incorporate side information, where A^D is the node attribute adjacency matrix that defines the pairwise attribute similarity between nodes and can be constructed according to custom user preferences (e.g. attributes with categorical or continuous values). VEGAS maintains flow rate maximization, performs localized summarization, and incorporates rich side information. In more detail, as shown in Figure 13, first the maximal influence graph G is computed from the input influence graph I by a rooted graph search that follows the standard BFS/DFS implementation from source node f. Then several matrices from the graph are generated: the topology similarity matrix M^G, and the optional matrices A^D and M^D. Finally non-negative matrix factorization (NMF) is used to solve the above optimization to obtain the community membership matrix H. Nodes are assigned to clusters by finding the maximum value in each row of H. Summaries are generated

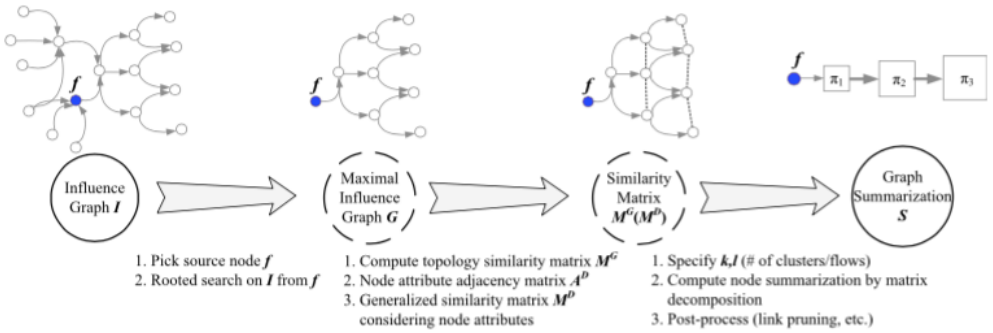

Fig. 13: *Algorithm pipeline of VEGAS [Shi, Tong, Tang, and Lin Shi et al.].*

after link pruning, which is performed to select the l best flows (links) for the final summary, dropping all other links between node clusters.

3.4. Pattern-mining-based methods

Pattern mining techniques aim at summarizing an input network via structural patterns that are also characterized by coherent node labels. For labeled graphs, the existing approaches attempt to identify frequent patterns, which are collapsed to supernodes, reducing the size of the input graph and increasing query efficiency.

SUBDUE [Cook and Holder 1994], which we already described as a compression-based approach in Section 3.2, falls into this category as well since it relies on finding frequent substructures and arranging them into supernodes. Similar in spirit, Chen et al. [2009] extracts frequently occurring subgraphs in heterogeneous graphs where nodes and edges are labeled (e.g. in a dependence graph, the vertices are labeled with program operations and the edges represent dependency relationships between them). The algorithm first generates partitions created by sampling nodes of the same label, resulting in multiple groups with consistent labels within. The partitioning/summarization is followed by frequent subgraph mining and verification (removal of false positives); these steps are performed in multiple iterations to find a lower bound on the false negative rate of frequent subgraph detection.

Query-oriented summarization [Čebirić et al. 2015] focuses on extracting the connection patterns followed by graphs in the Resource Description Framework (RDF) model, which are associated with the SPARQL query language. RDF data are stored as triples of `<subject, property, property value>` to present relationships. For example, the triple `<Harry Potter, HasAuthor, J.K. Rowling>` shows the relationship between the author and her work. The proposed algorithm summarizes the structure of the graph by generating a smaller graph that describes the types of nodes that occur at either end of the property edges. It groups together nodes and edges that follow the same connection format and attribute relationship. For example, consider two triples, `<Harry Potter, HasAuthor, J.K. Rowling>` and `<Harry Potter, HasPublisher, Bloomsbury>`: the algorithm will put `Source(HasAuthor)` and `Source(HasPublisher)` in the same node, as they are properties which share the same source type. In particular, for query-oriented summarization, the generated graph should support query-based representativeness and accuracy. For the former, queries with results on the original graph should have results on the summary graph; for the latter, a query that can be matched on the summary should also be matched on the original graph.

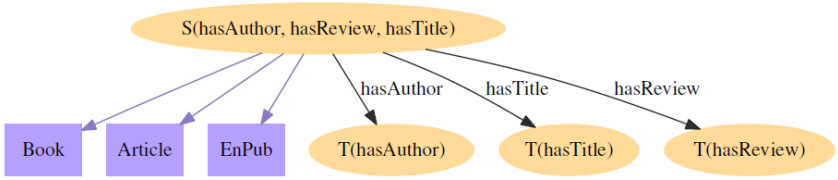

Fig. 14: *An example summary of the RDF graph (figure from [Čebirić et al. 2015])*

4. DYNAMIC GRAPH SUMMARIZATION: PLAIN NETWORKS

Analyzing large and complex data is challenging by itself, so adding the dimension of time makes the analysis even more challenging and time-consuming. Despite the challenges, the scenario is realistic: networks change over time – e.g. communication patterns with others via phone or social networks, the connection between servers in a network, the flow of information, news and rumors, the distance between connected vehicles, the information transmitted between devices in a smart home environment.

Table II: *Qualitative comparison of static graph summarization techniques. The first six columns describe the type of the input graph (e.g. with weighted/directed edges, and one/multiple types of node entities), followed by three algorithm-specific properties (i.e., user parameters, algorithmic compexity—linear on the number of edges or higher—, and type of output). The last column gives the final purpose of each approach. Notation: (1) * indicates that the algorithm can be extended to handle the corresponding type of input, but the authors do not provide details in the paper; (2) + means that at least one parameter can be set by the user, but it is not required (i.e., the algorithm provides a default value); (3) ✔(i) refers to interactive visualization where users can customize the summary or tune the resolution of final summaries.*

	Input Graph						Algorithmic Properties				
Method	**Unweighted**	**Weighted**	**Undirected**	**Directed**	**Homogeneous**	**Heterogeneous**	**Parameter-free**	**Linear**	**Core Technique**	**Output**	**Objective**
S-Node [Raghavan and Garcia-Molina 2003]	✔	✘	✘	✔	✔	✘	✔	✘	grouping	supergraph	query efficiency
SNAP/k-SNAP [Tian et al. 2008]	✔	✘	✔	✔*	✔(+)	✘	✔	✔	grouping	supergraph	query efficiency
CANAL [Zhang et al. 2010]	✔	✘	✔	✔*	✘	✘	✔	✘	grouping	supergraph	patterns
Probabilistic [Hassanlou et al. 2013]	✘	✔	✘	✔	✔	✘	✔	✔	grouping	supergraph	compression
ZKP [Shoaran et al. 2013]	✔	✘	✘	✔	✔	✘	✔	✔	grouping	supergraph	privacy
SUBDUE [Cook and Holder 1994]	✔	✘	✔	✔	✘	✔	✔	✘	compression	supergraph	patterns
AGSUMMARY [Wu et al. 2014]	✔	✘	✘	✔	✔	✘	✔	✔	compression	supergraph	compression
LSH-based [Khan et al. 2014]	✔	✘	✘	✔	✔	✘	✘	✔	compression	supergraph	compression
VEGAS [Shi, Tong, Tang, and Lin Shi et al.]	✔	✔*	✘	✔	✔	✘	✘	✔*	influence	supergraph	influence
Randomized [Chen et al. 2009]	✔	✘	✔	✘	✘	✔	✔	✘	pattern-mining	supergraph	patterns
RDF [Čebirić et al. 2015]	✔	✘	✘	✔	✘	✔	✔	✘	pattern-mining	connection patterns	visualization

For this reason, the temporal graph mining literature is rich, mostly focusing on: laws and patterns of graph evolution in [Ferlez et al. 2008b; Leskovec et al. 2008; Leskovec and Faloutsos 2007; Leskovec et al. 2005; Sun et al. 2008] and a comprehensive survey by Aggarwal and Subbian in [2014]; anomaly and change detection in streaming graphs [Aggarwal and Philip 2005] or time-evolving networks [Ferlez et al. 2008b; Koutra et al. 2015, 2013]; discovery of dense temporal cliques and bipartite cores using PARAFAC tensor decomposition and MDL ([Araujo et al. 2014; Sun et al. 2007],[Koutra et al. 2012]); mining of cross-graph quasi-cliques [Pei et al. 2005]; clustering using incremental static clustering [Xu et al. 2011] or a probabilistic approach based on mixed-membership blockmodels [Fu et al. 2009]; sampling of streaming graphs [Ahmed et al. 2013] and role discovery [Henderson et al. 2012; Rossi et al. 2012].

In this section, we focus on methods that summarize time-evolving networks. Summarization techniques for time-evolving networks have not been studied to the same extent as those for static networks, possibly because of the new challenges introduced by the dimension of time. The methods are sensitive to time granularity, or the timestamp of each network in the input graph sequence, which is often chosen arbitrarily (e.g. the granularity can be set to minutes, hours, days, weeks, months, years, or some other unit that makes sense in a given setting). The continuous, and sometimes irregular, change of real-world graphs also makes tracking evolution, defining an online "interestingness" measure, and visualizing them challenging. Nonetheless, dynamic graphs are ubiquitous. Summarizing or simplifying these graphs contributes to storage savings, reduction of data noise, and efficient analysis and visualization.

The problem of dynamic graph summarization can be described as:

PROBLEM 3.[Summarization of Dynamic, Plain Graphs]
Given a dynamic graph, which is observed as a set of streaming edges, or a sequence of adjacency matrices $\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_t$ corresponding to the static graphs $G_1, G_2, \ldots, G_t$
Find a temporal summary graph or a set of possibly overlapping temporal structures to **concisely describe** the given dynamic graph.

The summary is a time-evolving supergraph with supernodes and superedges, or else a sequence of sparsified graphs with fewer nodes/edges than the input dynamic graph.

A straightforward approach is to treat a time-evolving graph as a series of static graph snapshots and apply static graph summarization techniques. However, the analysis heavily depends on user-specified aggregation operations and the time granularity [Soundarajan et al. 2016], and there is no globally accepted way of picking the "right" time unit for the snapshots. By picking a very small time unit, the data to be independently processed increase significantly. On the other hand, a large time unit may lead to missing interesting dynamics. Moreover, processes can be quite unpredictable (e.g. bursty behavior) and can change abruptly over time. Adjusting the time unit of analysis may be the key to understanding and capturing the important dynamics.

An alternative way is to create an aggregate graph that summarizes the input dynamic network based on the recency and frequency of interactions, shown in Figure 15. This is referred to as an approximation graph in the literature [Cortes et al. 2001; Hill et al. 2006; Sharan and Neville 2008]. Specifically, the interactions between nodes in the approximation graph are aggregated over time in a weighted manner by applying kernel smoothing (e.g. exponential, inverse linear, linear, uniform). Intuitively, recent edges carry a larger weight than older edges in the graph approximation. A local and global threshold process can further simplify the graph approximation by pruning interactions of a small weight. The approximation graph has been shown to be useful for telecommunications fraud detection [Cortes et al. 2001], anomaly detection and prediction of user behavior in web logs and email networks [Hill et al. 2006], as well as attribute classification via relational classifier models [Sharan and Neville 2008].

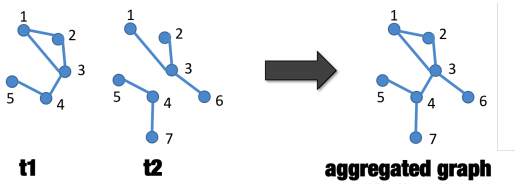

Fig. 15: *Example: Aggregated graph from time t1 and t2.*

The approximation graph can be used as input to any of the static graph summarization algorithms that we present in Section 2. However, this approach has the same shortcoming as the straightforward approach—namely, it depends on the time granularity of the input graph sequence. Probabilistic relational models (PRM) and relational Markov decision processes (RMDP, which are a sequence of PRMs forming a chain that follows a first-order Markov assumption) have also been used to model dynamic graphs [Guestrin et al. 2003], but they cannot model time-varying edges and treat them as fixed over time.

4.1. Compression-based methods

The techniques in this category use compression as a means of extracting meaningful patterns from temporal data. The only representative in this category is TIMECRUNCH [Shah et al. 2015], which succinctly describes a large, dynamic graph with a set of important temporal structures. Extending VOG [Koutra et al. 2014a], a static graph summarization presented in Section 2.3, the authors formalize the temporal graph summarization problem as an information-theoretic optimization problem where the goal is to identify the temporal behaviors

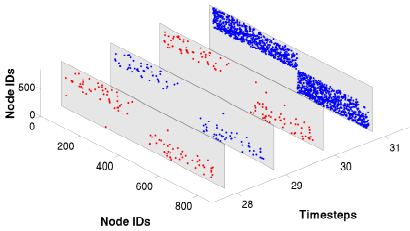

Fig. 16: TIMECRUNCH *[Shah et al. 2015] discovers meaningful structures in time-evolving graphs: Callers in a who-calls-whom network forming a one-shot near a bipartite core appearing strongly on Dec. 31.*

of local static structures that collectively minimize the global description length of the dynamic graph. A lexicon that describes various types of temporal behavior (e.g. flickering, periodic, one-shot) is introduced to augment the vocabulary of static graphs (e.g. stars, cliques, bipartite cores, chains). TIMECRUNCH (i) first identifies static structures in each timestamp, (ii) labels them using the static lexicon, (iii) stitches them together to find temporal structures, (iv) then labels those using the temporal lexicon, and (v) selects for the summary the temporal structures that help minimize the MDL cost of describing the time-evolving graph. Stitching static structures corresponds to evolution tracking, which is handled via iterative rank-1 singular value decomposition (SVD) to find potentially temporally-coherent structures, and then cosine similarity to ensure the temporal coherence of the discovered structures (i.e., there is sufficient similarity between static structures that comprise an evolution pattern). Figure 16 presents one example of a temporal pattern identified by TIMECRUNCH in real phonecall dataset.

4.2. Influence-based methods

Influence and diffusion processes are inherently time-evolving. The methods in this category aim at summarizing the influence mainly in social networks.

In Section 2.4, we present two techniques, CSI [Mehmood et al. 2013] and SPINE [Mathioudakis et al. 2011], which summarize social graphs by leveraging the information propagation and social influence processes. These approaches have a temporal aspect since they are summarizing inherently temporal activities in networks, but they operate on static graphs, where the directed edges capture influence.

Here we focus on OSNet [Qu et al. 2014], which summarizes interestingness-driven diffusion processes in dynamic graphs. The input is a stream of time-ordered interactions, represented as undirected edges between labeled nodes. The goal of OSNet is to capture cascades (i.e., diffusion process, such as spread of news) in a directed graph that reveals the flow of dynamics. The output summary consists of subgraphs with 'interesting' nodes from the original input graph, where interestingness is defined as a linear combination of the out-degree of a node (i.e., the number of nodes that it infects during the diffusion process), and the maximum 'propagation radius' (i.e., the length of the path from the root of the diffusion process to the node). The key ideas of the proposed algorithm, OSNet, are (i) to construct spreading trees and (ii) to compute the interestingness of a summary via its entropy and a threshold that can lead to fast convergence. OSNet outperforms static-based summarization techniques [Navlakha et al. 2008a; Toivonen et al. 2011] that give a summary per timestamp, since they are not suited for capturing the temporal dynamics; the former method depends on user-defined parameters, and the latter gives summaries with many disconnected cliques. VEGAS [Shi, Tong, Tang, and Lin Shi et al.], which we describe in Section 3.2, performs graph summarization also by maximizing the influence propagation, but only on static graphs.

Relatedly, Lin et al. [2008] focus on understanding the collective activity in a social group over time. Towards that goal, the authors propose to extract activity themes over time using non-negative matrix factorization (NMF) on a multi-graph (user-photo, user-comment, photo-tag, and comment-tag graphs) in order to get latent spaces for users and concepts. The top-k users and terms in the latent space define the 'important' actions which correspond to activity themes. Evolution of themes over time is tracked by applying cosine similarity between their corresponding latent spaces. We note that this bears similarities to the evolution tracking component of TIMECRUNCH [Shah et al. 2015], where cosine similarity was used to ensure temporal coherence. The themes are visualized as bubbles connected by edges whose length is proportional to the dissimilarity of the themes.

Connection to graph clustering, sparsification, and compression. As in the case of static graphs, techniques such as clustering, sparsification and compression can be cast as summarization methods for dynamic graphs. The techniques vary widely in terms of their assumptions: addition and/or deletion of nodes and/or edges over time. Some clustering methods extend heuristics that have been used for static graphs, such as modularity [Görke et al. 2010] or minimum-cut trees [Saha and Mitra 2007]. As we mention in Section 2, graph sketches [Ahn et al. 2012; Liberty 2013] are an alternative way of summarizing large amounts of data by applying linear projections. The property of linearity is fundamental as it makes sketches applicable to the analysis of streaming graphs in centralized or even distributed settings, where they are partitioned in multiple servers in MapReduce [Dean and Ghemawat 2004] or a different environment. One-pass and other efficient streaming algorithms with their theoretical analysis are given in [Ahn et al. 2012].

Work on compressing dynamic graphs for storage includes lossy compression of time-evolving graphs [Henecka and Roughan 2015], and encoding of a dynamic, weighted graph as a three-dimensional array (tensor) by reducing its heterogeneity and guaranteeing that the compression error is within bounds [Liu et al. 2012]. The method is based on hierarchical clusters of edge weights and graph compression using run-length encoding by traversing along the time dimension first, and then along the vertex dimensions of the tensor. A key property of the proposed method is that it maintains the connectivity of the graph as defined by the average shortest paths over all pairs of connected nodes, and thus can handle such queries with good approximations.

Table III: *Qualitative comparison of dynamic graph summarization techniques. The first six columns describe the type of the input graph (e.g. with weighted/directed edges, and one/multiple types of node entities), followed by three algorithm-specific properties (i.e., user parameters, algorithmic compexity—linear on the number of edges or higher—, and type of output). The last column gives the final purpose of each approach.*

	Input Graph Properties						Algorithmic Properties				
Method	**Unweighted**	**Weighted**	**Undirected**	**Directed**	**Homogeneous**	**Heterogeneous**	**Parameter-free**	**Linear**	**Core Technique**	**Output**	**Objective**
TIMECRUNCH [Shah et al. 2015]	✔	✘	✔	✘	✔	✔	✘	✔	compression	ranked list of temporal structures	temporal patterns, visualization
OSNET [Qu et al. 2014]	✔	✘	✔	✘	✔	✘	✘	✔	influence	subgraphs of diffusion over time	influnence
Social Activity [Lin et al. 2008]	✔	✘	✔	✘	✘	✔	✘	✘	influence	themes over time	influence, visualization

5. GRAPH SUMMARIZATION IN REAL-WORLD APPLICATIONS

As we mention in the introduction, summarization helps mitigate information overload. In this section, we discuss real-world applications of graph summarization, which are myriad, diverse, and relevant to many domains.

5.1. Summarization for query efficiency

Graph summarization can greatly improve graph query efficiency. Depending on the query type and the structure of the summary, some queries can easily be answered by the summary with good accuracy. For example, in the case of a social network graph, such a query might be whether an edge exists between two nodes; this query can be answered on a space-efficient summary of an expected adjacency matrix [LeFevre and Terzi 2010; Riondato et al. 2014].

One important application of improving query efficiency is on Web graphs [Boldi and Vigna 2004b; Raghavan and Garcia-Molina 2003]. In the latter work, a set of complex queries executed over the Stanford WebBase repository measure the performance of their S-Node (Section 3.1) representations for local graph navigation. These queries (find out-neighborhood of pages; count number of links; get union of out-neighborhood and in-neighborhood of pages, as well as in-neighborhood for different pages; compute induced graphs; obtain intersection of out-neighborhoods of different pages) represent user preference and illustrate a variety of graph navigation operations that could be employed in queries. The S-Node representation outperforms other representation schemes (relational database, connectivity server, and files of adjacency lists), reducing navigation time by an order of magnitude. This is because the S-Node scheme loads only a relatively small number of intranode and superedge graphs, hence there are savings in disk I/Os.

Other examples of improved query efficiency are estimating edge weights, counting the number of triangles [Riondato et al. 2014], and ranking and finding communities [Kang et al. 2011].

5.2. Summarization for visualization

Summarization is an essential tool for visualizing graph data that would otherwise be too large to display. OntoVis [Shen et al. 2006] (Section 2.2) applies on a large heterogeneous movie network. The graph has 8 node types (person, movie, role etc.) with 35,000 nodes and 108,000 links; even a graph of this size is impossible to fit on a desktop screen. To investigate the relationships between persons and roles, the authors observe the derived abstraction and point out a role-actor relationship where a good actor should be able to play different roles (e.g. actors like Woody Allen and Sandra Bullock, who play three different types of roles).

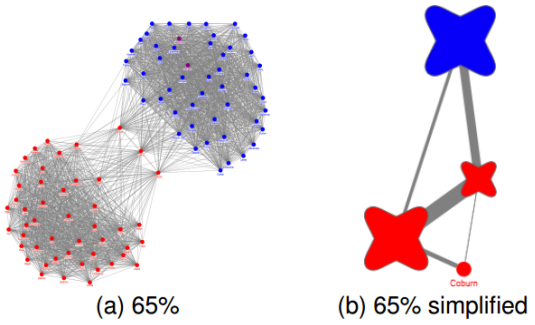

Fig. 17: *US Senate co-voting network of 100 nodes (red nodes are Republican senators, blue are Democratic) simplified using clique motif glyphs. The percent similarity in voting patterns is an attribute of each edge, where only similarity values of 65% or greater are considered [Dunne and Shneiderman 2013].*

Other works that perform visualization on top of summarization include VoG [Koutra et al. 2014b], which visualizes structures of specific types (e.g. cliques, bipartite cores), and Motif Simplification [Dunne and Shneiderman 2013], which visualizes simplified networks of up to 8,000 nodes with glyphs, an example of which is given in Figure 17.

5.3. Summarization for compression

Compression and summarization are closely related; compression may serve as the means to summarization, and vice-versa.

MDL is a crucial and widely used principle for compression. Navlakha et al. [2008a] (Section 2.3) are among the first to compute graph summaries using MDL; the authors test their algorithms for both exact and approximate MDL representations on datasets of various domains and sizes, such as a web graph, an autonomous system network, and a lexical database of English words. Indeed, visualization and trend analysis on the small and less cluttered supergraph summary (about 10% of the original graph) shows many interesting patterns that are not visible in the original graph. For example,

Figure 18 shows a bipartite core with a large set of nodes connecting to a single other node, which would be hard to identify in the original graph but stands out in the summary because these nodes are condensed into a single large supernode.

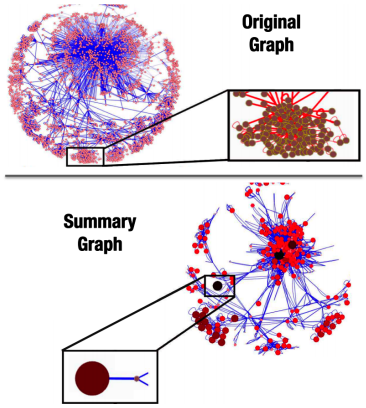

Fig. 18: *The original graph contains a hard-to-identify bipartite core (magnified at the top). In the summary, it has been condensed into a supernode, making it easier to identify the relationship between the original nodes [Navlakha et al. 2008b].*

Another example is VoG [Koutra et al. 2014b], which is applied to multiple collaboration graphs, including Wikipedia co-editor graphs. VoG reveals edit wars in controversial topics, which appear as near bipartite cores. More examples include the compression of various web graphs by VNM [Buehrer and Chellapilla 2008], the Eulerian data structure that compresess social networks while preserving efficiency for neighbor queries [Maserrat and Pei 2010], and community-preserving lossy compression [Maserrat and Pei 2012], which shows effectiveness in terms of both reachability and pattern queries on attributed graphs such as a Youtube video graph, a California Web graph, and a citation graph.

5.4. Summarization for pattern discovery

Nodes and edges frequently form patterns that are not immediately noticeable, but summarization can effectively extract these patterns. Pattern discovery is often used to determine the behavior of elements in a graph based on their structure. Consider the `Wikipedia-Controversy` dataset, where nodes are Wikipedia contributors and edges connect users who edit the same part of the article. VoG [Koutra et al. 2014b] applied on this graph for the purpose of extracting the top 10 most informative structures yields 8 stars and 2 bipartite graphs. The centers of the stars correspond to Wikipedia admins or heavy contributors; the bipartite cores correspond to edit wars between two groups of users, such as vandals and responsible editors, on a controversial topic.

SUBDUE [Cook and Holder 1994], one of the most famous frequent-pattern mining methods, has applications as diverse as chemical compound analysis, scene analysis, and CAD circuit design analysis. For example, SUBDUE discovers substructures in chemical compound graphs where labeled vertices are individual atoms and labeled edges are bonds; in particular, the building-block components that are heavily used, such as isoprene units for rubber compounds, are discovered.

5.5. Summarization for influence extraction

Influence extraction is a long-standing research focus of graph mining. Accordingly, contribution to this goal has been made from many angles.

For example, [Li and Lin 2009] tries to assist human decision-making via egocentric abstraction. The authors use a simulated heterogeneous crime dataset with nodes as gangs and edges as relations between them. It is demonstrated that presenting users with the abstraction of the graph can lead to more accurate, efficient, and confident identification of high-level crime-committing gangs. Furthermore, it is demonstrated that each abstraction view captures different parts of key evidence to some extent, such as “the gang has hired some middleman intending to pursue something illegal”.

Other examples include COARSENET [Purohit et al. 2014]: applied on cascade network Flixster of 56,000 nodes and 560,000 edges, it is demonstrated that a very large fraction of movies propagate in a small number of groups with a multi-modal distribu-

tion, suggesting movies have multiple scales of spread. In the community-level social influence analysis on Yahoo! and Twitter graphs by CSI [Mehmood et al. 2013], the authors register almost no correlation between influence and link probabilities. In other words, influence relationships do not in general exhibit any clear structure, and even if a community is dense, it does not necessarily exhibit strong internal influence.

6. CONCLUSION

In this survey we presented the state-of-the-art in graph summarization. Based on the type of input graph and the main technical methodology, we proposed a taxonomy and categorized existing graph summarization algorithms. For each algorithm, we provided an introduction to its key idea and explored the relations between relevant works. Furthermore, we drew connections between graph summarization to other data representations and provided examples of real-life applications for each category of algorithms. In the following we point readers to some important open problems in the field of graph summarization.

6.1. Open Research Problems

While research in the field of graph summarization has recently picked up, the field is still new and in need of exploration.

With regard to input data, there does not yet exist work on summarizing temporal graphs with side information. Many real-world domains like social networks can easily be modeled as temporal attributed graphs, and summarization for such graphs has the potential to capture changing patterns of structure *and* attribute properties over time.

Even foregoing the temporal aspect, summarization of *various graph types* has not been thoroughly explored. An example is the multi-layer graph, which is an important model for real-world data such as Web graphs [Laura et al. 2002]. Another example is multiview data, where the graph can be analyzed with respect to different edge types. For instance, Twitter data can include edges for follow, retweet, and messaging, each of which can have its own adjacency matrix. Spatiotemporal graphs, which have applications in map representations by tracking the trajectories of people over time, can be used with summarization to extract commonly traveled paths and allow for personalizable mapping. Furthermore, as data become increasingly richer, we will need to understand graphs that are incorporated with other types of data sources, e.g., time series associated with nodes. Collective summarization of this type of data or multiple graphs with dependencies other than temporal is to be studied.

A lack of standard *evaluation techniques* in graph summarization leads to non-objective, application-based evaluation. If the method applies to compression, the evaluation method becomes compression quality, whereas if the method applies to query efficiency, the evaluation method becomes query latency. Some common evaluation metrics can render comparing new methods with the state-of-the-art approaches easier. For example, metrics that evaluate supergraphs base quality on sparsity, least information loss, and ease of visualization.

Graph summarization is also important in *time-sensitive settings*, but to our knowledge this has not yet been addressed in the literature. Real-time graph summarization is an important application that can be done in a streaming setting. The usefulness of stored summaries also decreases as a graph changes over time. This can be countered by designing methods that allow for incremental changes to summaries. Incremental summarization keeps the summaries up to date and spares burden of generating new summaries from scratch.

As for query efficiency-oriented graph summarization, numerous methods that address query efficiency on graph summaries exist, but they either only offer approximate

querying or are limited to very specific queries. There is still work to be done on lossless compression with *general-purpose queries*.

An aspect that is often overlooked is that of *automating the process of understanding* the generated summaries with further analysis. For this purpose, natural language processing techniques can be employed for analysis of side information such as attributes. Moreover, although in labeled graphs all methods look to group nodes with cohesive attribute values, there are many applications in which heterogeneous clusters are crucial. For example, such clusters could contribute to anomaly detection, or they could refer to groups with desired diversity (e.g. in an academic setting).

Finally, a new, promising direction is graph summarization based on *deep node representations* that are learned automatically from the context encoded in the graph. Learning node representations is a new area in data mining that has attracted significant interest in the last couple of years [Grover and Leskovec 2016; Perozzi et al. 2014; Tang et al. 2015; Wang et al. 2016]. There are already summarization methods that rely on latent node representations (e.g., via factorization) or node/egonet/k-hop neighborhood features that are manually selected for extraction. Given the success of deep learning in other tasks, studying the potential of deep node representations in summarization and compression of interconnected data seems promising.

Overall, graph summarization methods are becoming increasingly important and useful as the volume of available interconnected data grows rapidly. While we overview several formulations of graph summarization already studied, we conclude by noting that many promising directions in the field remain unexplored and thus full of potential for impact.

References

Charu Aggarwal and Karthik Subbian. 2014. Evolutionary Network Analysis: A Survey. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 10:1–10:36.

Charu C Aggarwal. 2015. *Data Mining: The Textbook*. Springer.

Charu C Aggarwal and S Yu Philip. 2005. Online Analysis of Community Evolution in Data Streams. In *Proceedings of the 5th SIAM International Conference on Data Mining (SDM), Newport Beach, CA*. SIAM.

Charu C Aggarwal and Haixun Wang. 2010. A Survey of Clustering Algorithms for Graph Data. In *Managing and mining graph data*. Springer, 275–301.

Amr Ahmed, Nino Shervashidze, Shravan Narayanamurthy, Vanja Josifovski, and Alexander J Smola. 2013. Distributed Large-Scale Natural Graph Factorization. In *Proceedings of the 22nd International Conference on World Wide Web (WWW), Rio de Janeiro, Brazil*. International World Wide Web Conferences Steering Committee, 37–48.

Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. 2012. Graph Sketches: Sparsification, Spanners, and Subgraphs. In *Proceedings of the 31st ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. 5–14.

Alfred V. Aho, M. R. Garey, and Jeffrey D. Ullman. 1972. The Transitive Reduction of a Directed Graph. *Siam Journal on Computing* 1 (1972), 131–137. Issue 2.

Leman Akoglu*, Duen Horng Chau*, U Kang*, Danai Koutra*, and Christos Faloutsos. 2012. OPAvion: Mining and Visualization in Large Graphs. In *Proceedings of the 2012 ACM International Conference on Management of Data (SIGMOD), Scottsdale, AZ*. ACM, 717–720.

Leman Akoglu, Mary McGlohon, and Christos Faloutsos. 2010. OddBall: Spotting Anomalies in Weighted Graphs. In *Proceedings of the 14th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), Hyderabad, India*.

Leman Akoglu, Hanghang Tong, Jilles Vreeken, and Christos Faloutsos. 2012. Fast

and Reliable Anomaly Detection in Categorical Data. In *Proceedings of the 21st ACM Conference on Information and Knowledge Management (CIKM), Maui, Hawaii*. ACM.

Leman Akoglu, Jilles Vreeken, Hanghang Tong, Nikolaj Tatti, and Christos Faloutsos. 2013. Mining Connection Pathways for Marked Nodes in Large Graphs. In *Proceedings of the 13th SIAM International Conference on Data Mining (SDM), Texas-Austin, TX*. SIAM.

Charles J Alpert, Andrew B Kahng, and So-Zen Yao. 1999. Spectral partitioning with multiple eigenvectors. *Discrete Applied Mathematics* 90, 1 (1999), 3–26.

Alberto Apostolico and Guido Drovandi. 2009. Graph Compression by BFS. *Algorithms* 2, 3 (2009), 1031–1044.

Miguel Araujo, Spiros Papadimitriou, Stephan Günnemann, Christos Faloutsos, Prithwish Basu, Ananthram Swami, Evangelos E Papalexakis, and Danai Koutra. 2014. Com2: Fast Automatic Discovery of Temporal ("Comet") Communities. In *Proceedings of the 18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), Tainan, Taiwan*. Springer, 271–283.

Joshua D. Batson, Daniel A. Spielman, Nikhil Srivastava, and Shang-Hua Teng. 2013. Spectral Sparsification of Graphs: Theory and Algorithms. *Commun. ACM* 56, 8 (2013), 87–94.

Enrico Bertini and Giuseppe Santucci. 2004. By Chance is not Enough: Preserving Relative Density through non Uniform Sampling. In *Proceedings of the Information Visualisation*.

Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast Unfolding of Communities in Large Networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008.

Paolo Boldi and Sebastiano Vigna. 2004a. The Webgraph Framework I: Compression Techniques. In *Proceedings of the 13th International Conference on World Wide Web (WWW), New York, NY*.

Paolo Boldi and Sebastiano Vigna. 2004b. The WebGraph Framework I: Compression techniques. In *International World Wide Web Conference*. 595–602.

Gregory Buehrer and Kumar Chellapilla. 2008. A Scalable Pattern Mining Approach to Web Graph Compression With Communities. In *Proceeding of the 1st ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, 95–106.

Gemma Casas-Garriga. 2005. Summarizing Sequential Data with Closed Partial Orders. In *Proceedings of the 5th SIAM International Conference on Data Mining (SDM), Newport Beach, CA*. 380–391.

Šelja Čebirić, François Goasdoué, and Ioana Manolescu. 2015. Query-Oriented Summarization of RDF Graphs. *Proceedings of the VLDB Endowment* 8, 12 (2015).

Deepayan Chakrabarti, Spiros Papadimitriou, Dharmendra S. Modha, and Christos Faloutsos. 2004. Fully Automatic Cross-associations. In *Proceedings of the 10th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Seattle, WA*. 79–88.

Varun Chandola and Vipin Kumar. 2005. Summarization – Compressing Data Into an Informative Representation. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM), Houston, TX*. 98–105.

Duen Horng Chau, Aniket Kittur, Jason I. Hong, and Christos Faloutsos. 2011. Apolo: Making Sense of Large Network Data by Combining Rich User Interaction and Machine Learning. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*.

Chen Chen, Cindy X Lin, Matt Fredrikson, Mihai Christodorescu, Xifeng Yan, and Jiawei Han. 2009. Mining Graph Patterns Efficiently Via Randomized Summaries.

Proceedings of the 35nd International Conference on Very Large Data Bases (VLDB), Lyon, France 2, 1 (2009), 742–753.

Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, Michael Mitzenmacher, Alessandro Panconesi, and Prabhakar Raghavan. 2009. On Compressing Social Networks. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Paris, France*. 219–228.

Yongwook Choi and Wojciech Szpankowski. 2012. Compression of Graphical Structures: Fundamental Limits, Algorithms, and Experiments. *IEEE Transactions on Information Theory* 58, 2 (2012), 620–638.

Rudi Cilibrasi and Paul Vitányi. 2005. Clustering by Compression. *IEEE Transactions on Information Technology* 51, 4 (2005), 1523–1545.

Diane J. Cook and Lawrence B. Holder. 1994. Substructure Discovery Using Minimum Description Length and Background Knowledge. *Journal of Artificial Intelligence Research* 1 (1994), 231–255.

Robson L. F. Cordeiro, Fan Guo, Donna S. Haverkamp, James H. Horne, Ellen K. Hughes, Gunhee Kim, Agma J. M. Traina, Caetano Traina Jr., and Christos Faloutsos. 2010. QMAS: Querying, Mining And Summarization of Multi-modal Databases. *ICDM* (2010), 785–790.

Graham Cormode and S. Muthukrishnan. 2005. Summarizing and Mining Skewed Data Streams.. In *SDM*.

Corinna Cortes, Daryl Pregibon, and Chris Volinsky. 2001. Communities of Interest. In *Proceedings of the 4th International Conference on Advances in Intelligent Data Analysis*. Springer-Verlag, 105–114.

Uros Damnjanovic, Virginia Fernandez Arguedas, Ebroul Izquierdo, and José M. Martínez. 2008. Event Detection and Clustering for Surveillance Video Summarization.. In *9th International Workshop on Image Analysis for Multimedia Interactive Services*. IEEE Computer Society, 63–66.

Jeffrey Dean and Sanjay Ghemawat. 2004. MapReduce: Simplified Data Processing on Large Clusters. *Proceedings of the 6th Symposium on Operating Systems Design and Implementation, San Francisco, CA* (Dec. 2004).

Inderjit Dhillon, Yuqiang Guan, and Brian Kulis. 2005. A Fast Kernel-based Multilevel Algorithm for Graph Clustering. In *Proceedings of the 11th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Chicago, IL*. ACM, 629–634.

Inderjit S. Dhillon, Subramanyam Mallela, and Dharmendra S. Modha. 2003. Information-Theoretic Co-clustering. In *Proceedings of the 9th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Washington, DC*. 89–98.

Laxman Dhulipala, Igor Kabiljo, Brian Karrer, Giuseppe Ottaviano, Sergey Pupyrev, and Alon Shalita. 2016. Compressing Graphs and Indexes with Recursive Graph Bisection. *arXiv preprint arXiv:1602.08820* (2016).

Vincent Dubois and Cécile Bothorel. 2005. Transitive Reduction for Social Network Analysis and Visualization. In *2005 IEEE / WIC / ACM International Conference on Web Intelligence (WI)*. 128–131.

Cody Dunne and Ben Shneiderman. 2013. Motif Simplification: Improving Network Visualization Readability with Fan, Connector, and Clique Glyphs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 3247–3256.

Michael Elkin and David Peleg. 2005. Approximating k-spanner Problems for k ¿ 2. *Theoretical Computer Science* 337, 1 (2005), 249–277.

Wenfei Fan, Jianzhong Li, Xin Wang, and Yinghui Wu. 2012. Query Preserving Graph Compression. In *Proceedings of the 2012 ACM International Conference on Manage-*

ment of Data (SIGMOD), Scottsdale, AZ. ACM, 157–168.
Katherine Faust and Stanley Wasserman. 1992. Blockmodels: Interpretation and Evaluation. *Social Networks* 14, 1-2 (1992), 5–61.
Jing Feng, Xiao He, Nina Hubig, Christian Böhm, and Claudia Plant. 2013. Compression-Based Graph Mining Exploiting Structure Primitives. In *Proceedings of the 14th IEEE International Conference on Data Mining (ICDM), Dallas, Texas*. IEEE, 181–190.
Jure Ferlez, Christos Faloutsos, Jure Leskovec, Dunja Mladenic, and Marko Grobelnik. 2008a. Monitoring Network Evolution using MDL. In *Proceedings of the 24th International Conference on Data Engineering (ICDE), Cancun, Mexico*. 1328–1330.
Jure Ferlez, Christos Faloutsos, Jure Leskovec, Dunja Mladenic, and Marko Grobelnik. 2008b. Monitoring Network Evolution using MDL. In *Proceedings of the 24th International Conference on Data Engineering, ICDE 2008, April 7-12, 2008, Cancún, México*. 1328–1330.
Wenjie Fu, Le Song, and Eric P. Xing. 2009. Dynamic Mixed Membership Blockmodel for Evolving Networks. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML '09)*. ACM, New York, NY, USA, 329–336.
Johannes Gehrke, Edward Lui, and Rafael Pass. 2003. Towards Privacy for Social Networks: a Zero-Knowledge Based Definition of Privacy. In *Proceedings of the 8th Conference on Theory of Cryptography*. IACR, 432–449.
Mina Ghashami, Edo Liberty, and Jeff M Phillips. 2016. Efficient Frequent Directions Algorithm for Sparse Matrices. *arXiv preprint arXiv:1602.00412* (2016).
Robert Görke, Pascal Maillard, Christian Staudt, and Dorothea Wagner. 2010. Modularity-Driven Clustering of Dynamic Graphs. In *Experimental Algorithms, 9th International Symposium, SEA 2010, Ischia Island, Naples, Italy, May 20-22, 2010. Proceedings*. 436–448.
Szymon Grabowski and Wojciech Bieniecki. 2014. Tight and Simple Web Graph Compression For Forward and Reverse Neighbor Queries. *Discrete Applied Mathematics* 163 (2014), 298–306.
Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Francisco, CA*. ACM.
C. Guestrin, D. Koller, C. Gearhart, and N. Kanodia. 2003. Generalizing Plans to New Environments in Relational MDPs. In *Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*. Acapulco, Mexico.
Mohammad Al Hasan, Nesreen K. Ahmed, and Jennifer Neville. 2013. Network Sampling: Methods and Applications. https://www.cs.purdue.edu/homes/neville/courses/NetworkSampling-KDD13-final.pdf. (2013).
Nasrin Hassanlou, Maryam Shoaran, and Alex Thomo. 2013. Probabilistic Graph Summarization. In *Web-Age Information Management*. Springer, 545–556.
Keith Henderson, Brian Gallagher, Tina Eliassi-Rad, Hanghang Tong, Sugato Basu, Leman Akoglu, Danai Koutra, Christos Faloutsos, and Lei Li. 2012. RolX: Structural Role Extraction & Mining in Large Graphs. In *Proceedings of the 18th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Beijing, China*. ACM, 1231–1239.
W. Henecka and M. Roughan. 2015. Lossy Compression of Dynamic, Weighted Graphs. In *Future Internet of Things and Cloud (FiCloud), 2015 3rd International Conference on*. 427–434.
S. Hill, D. Agarwal, R. Bell, and C. Volinsky. 2006. Building an Effective Representation for Dynamic Networks. *Journal of Computational and Graphical Statistics* 15 (2006), 1–25.
Christian Hübler, Hans-Peter Kriegel, Karsten Borgwardt, and Zoubin Ghahramani.

2008. Metropolis Algorithms for Representative Subgraph Sampling. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining (ICDM '08)*. IEEE Computer Society, Washington, DC, USA, 283–292.

U Kang, Jay-Yoon Lee, Danai Koutra, and Christos Faloutsos. 2014. Net-Ray: Visualizing and Mining Web-Scale Graphs.. In *Proceedings of the 18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), Tainan, Taiwan*.

U Kang, Hanghang Tong, Jimeng Sun, Ching-Yung Lin, and Christos Faloutsos. 2011. Gbase: A Scalable and General Graph Management System. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*. ACM, 1091–1099.

U. Kang, Charalampos E. Tsourakakis, and Christos Faloutsos. 2009. PEGASUS: A Peta-Scale Graph Mining System - Implementation and Observations. *Proceedings of the 9th IEEE International Conference on Data Mining (ICDM), Miami, FL* (2009).

George Karypis and Vipin Kumar. 1999. Multilevel k-way Hypergraph Partitioning. In *Proceedings of the IEEE 36th Conference on Design Automation Conference (DAC), New Orleans, LA*. 343–348.

David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Conference of the ACM Special Interest Group on Knowledge Discovery and Data Mining*. ACM Press, New York, NY.

Kifayat Ullah Khan, Waqas Nawaz, and Young-Koo Lee. 2014. Set-Based Unified Approach for Attributed Graph Summarization. In *IEEE Fourth International Conference on Big Data and Cloud Computing (BdCloud)*. IEEE, 378–385.

Jerry Kiernan and Evimaria Terzi. 2009. EventSummarizer: a tool for summarizing large event sequences. In *Proceedings of the 12th International Conference on Extending Database Technology (EDBT), Saint-Petersburg, Russia*. 1136–1139.

Arne Koopman and Arno Siebes. 2008. Discovering Relational Items Sets Efficiently. In *Proceedings of the 8th SIAM International Conference on Data Mining (SDM), Atlanta, GA*. 108–119.

Arne Koopman and Arno Siebes. 2009. Characteristic Relational Patterns. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Paris, France*. 437–446.

Danai Koutra, Di Jin, Yuanchi Ning, and Christos Faloutsos. 2015. Perseus: An Interactive Large-Scale Graph Mining and Visualization Tool. In *VLDB (demo)*.

Danai Koutra, U Kang, Jilles Vreeken, and Christos Faloutsos. 2014a. VoG: Summarizing and Understanding Large Graphs. In *Proceedings of the 14th SIAM International Conference on Data Mining (SDM), Philadelphia, PA*. 91–99.

Danai Koutra, U Kang, Jilles Vreeken, and Christos Faloutsos. 2014b. VOG: Summarizing and Understanding Large Graphs. In *Proceedings of the 14th SIAM International Conference on Data Mining (SDM), Philadelphia, PA*. SIAM.

Danai Koutra, Evangelos E Papalexakis, and Christos Faloutsos. 2012. TensorSplat: Spotting Latent Anomalies in Time. In *Informatics (PCI), 2012 16th Panhellenic Conference on*. IEEE, 144–149.

Danai Koutra, Neil Shah, Joshua Vogelstein, Brian Gallagher, and Christos Faloutsos. 2015. DeltaCon: A Principled Massive-Graph Similarity Function with Attribution. *ACM Transactions on Knowledge Discovery from Data* (2015).

Danai Koutra, Joshua Vogelstein, and Christos Faloutsos. 2013. DeltaCon: A Principled Massive-Graph Similarity Function. In *Proceedings of the 13th SIAM International Conference on Data Mining (SDM), Texas-Austin, TX*. 162–170.

Luigi Laura, Stefano Leonardi, Guido Caldarelli, and Paolo De Los Rios. 2002. A Multi-Layer Model for the Web Graph. In *In On-line proceedings of the 2nd International Workshop on Web Dynamics*.

Matthijs van Leeuwen, Jilles Vreeken, and Arno Siebes. 2006. Compression Picks the Item Sets that Matter. In *Proceedings of the 10th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD), Berlin, Germany*. 585–592.

Kristen LeFevre and Evimaria Terzi. 2010. GraSS: Graph Structure Summarization. In *Proceedings of the 10th SIAM International Conference on Data Mining (SDM), Columbus, OH*. SIAM, 454–465.

Jure Leskovec, Lars Backstrom, Ravi Kumar, and Andrew Tomkins. 2008. Microscopic evolution of social networks. In *Proceedings of the 14th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Las Vegas, NV*. 462–470.

Jure Leskovec and Christos Faloutsos. 2006. Sampling from Large Graphs. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '06)*. ACM, New York, NY, USA, 631–636.

Jure Leskovec and Christos Faloutsos. 2007. Scalable modeling of real graphs using Kronecker multiplication. In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvalis, Oregon, USA, June 20-24, 2007*. 497–504.

Jure Leskovec, Jon M. Kleinberg, and Christos Faloutsos. 2005. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, Illinois, USA, August 21-24, 2005*. 177–187.

Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. 2014. *Mining of Massive Datasets*. Cambridge University Press.

Cheng-Te Li and Shou-De Lin. 2009. Egocentric Information Abstraction for heterogeneous social networks. In *International Conference on Advances in Social Network Analysis and Mining (ASONAM)*. IEEE, 255–260.

Edo Liberty. 2013. Simple and Deterministic Matrix Sketching. In *Proceedings of the 19th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Chicago, IL*. ACM, 581–588.

Yongsub Lim, U Kang, and Christos Faloutsos. 2014. SlashBurn: Graph Compression and Mining beyond Caveman Communities. *IEEE Transactions on Knowledge and Data Engineering* 26, 12 (2014), 3077–3089.

Shou-De Lin, Mi-Yen Yeh, and Cheng-Te Li. 2013. Sampling and Summarization for Social Networks. http://mslab.csie.ntu.edu.tw/tut-pakdd13/, *PAKDD 2013 Tutorial* (2013).

Xuemin Lin, Qing Liu, Yidong Yuan, and Xiaofang Zhou. 2003. Multiscale Histograms: Summarizing Topological Relations in Large Spatial Datasets. In *International Conference on Very Large Databases (VLDB)*, Johann Christoph Freytag, Peter C. Lockemann, Serge Abiteboul, Michael J. Carey, Patricia G. Selinger, and Andreas Heuer (Eds.). Morgan Kaufmann Publishers, Inc, Berlin, Germany, 814–825.

Yu-Ru Lin, Hari Sundaram, and Aisling Kelliher. 2008. Summarization of Social Activity Over Time: People, Actions and Concepts in Dynamic Networks. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (CIKM), Napa Valley, CA*. 1379–1380.

Bing Liu, Wynne Hsu, and Yiming Ma. 1999. Pruning and Summarizing the Discovered Associations. In *Fifth ACM SIGKDD Intl Conf. Knowledge Discovery and Data Mining (KDD 99)*. 145–154.

Chunyang Liu and Ling Chen. 2016. Summarizing Uncertain Transaction Databases by Probabilistic Tiles. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 4375–4382.

Wei Liu, Andrey Kan, Jeffrey Chan, James Bailey, Christopher Leckie, Jian Pei, and Ramamohanarao Kotagiri. 2012. On Compressing Weighted Time-evolving Graphs. In

Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM '12). ACM, New York, NY, USA, 2319–2322.

Yike Liu, Neil Shah, and Danai Koutra. 2015. An Empirical Comparison of the Summarization Power of Graph Clustering Methods. *arXiv preprint arXiv:1511.06820* (2015).

Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M. Hellerstein. 2012. Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud. *Proceedings of the VLDB Endowment* 5, 8 (2012), 716–727.

Antonio Maccioni and Daniel J. Abadi. 2016. Scalable Pattern Matching over Compressed Graphs via Dedensification. In *Proceedings of the 22nd ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Francisco, CA*. ACM, 1755–1764.

Arun S. Maiya and Tanya Y. Berger-Wolf. 2010. Sampling Community Structure. In *Proceedings of the 19th International Conference on World Wide Web (WWW), Raleigh, NC*. ACM, New York, NY, USA, 701–710.

Michael Mampaey, Nikolaj Tatti, and Jilles Vreeken. 2011a. Tell Me What I Need To Know: Succinctly Summarizing Data with Itemsets. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*. ACM, 573–581.

M. Mampaey, J. Vreeken, and N. Tatti. 2011b. *Summarizing Data with Itemsets using Maximum Entropy Models*. Technical Report 2011/02. University of Antwerp. http://adrem.ua.ac.be/publications.

Koji Maruhashi, Fan Guo, and Christos Faloutsos. 2011. MultiAspectForensics: Pattern Mining on Large-Scale Heterogeneous Networks with Tensor Analysis. In *Proceedings of the 2011 International Conference on Advances in Social Networks Analysis and Mining*. 203–210.

Hossein Maserrat and Jian Pei. 2010. Neighbor Query Friendly Compression of Social Networks. In *Proceedings of the 16th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Washington, DC*.

Hossein Maserrat and Jian Pei. 2012. Community Preserving Lossy Compression of Social Networks. In *Proceedings of the 12th IEEE International Conference on Data Mining (ICDM), Brussels, Belgium*. IEEE, 509–518.

Michael Mathioudakis, Francesco Bonchi, Carlos Castillo, Aristides Gionis, and Antti Ukkonen. 2011. Sparsification of Influence Networks. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*. 529–537.

Yasir Mehmood, Nicola Barbieri, Francesco Bonchi, and Antti Ukkonen. 2013. Csi: Community-Level Social Influence Analysis. In *Machine Learning and Knowledge Discovery in Databases*. Springer, 48–63.

Pauli Miettinen and Jilles Vreeken. 2011. Model Order Selection for Boolean Matrix Factorization. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*. 51–59.

Pauli Miettinen and Jilles Vreeken. 2014. MDL4BMF: Minimum Description Length for Boolean Matrix Factorization. *ACM Transactions on Knowledge Discovery from Data* 8, 4 (2014), 1–30.

Saket Navlakha, Rajeev Rastogi, and Nisheeth Shrivastava. 2008a. Graph Summarization with Bounded Error. In *Proceedings of the 2008 ACM International Conference on Management of Data (SIGMOD), Vancouver, BC*. 419–432.

Saket Navlakha, Rajeev Rastogi, and Nisheeth Shrivastava. 2008b. Graph Summarization With Bounded Error. In *Proceedings of the 2008 ACM International Conference on Management of Data (SIGMOD), Vancouver, BC*. ACM, 419–432.

Mark E. J. Newman and Michelle Girvan. 2004. Finding and Evaluating Community Structure in Networks. *Physical Review E* 69, 2 (Feb. 2004), 026113+.

Carlos Ordonez, Norberto Ezquerra, and Cesar A. Santana. 2006. Constraining and summarizing association rules in medical data. *Knowledge and Information Systems* 9, 3 (2006), 259–283.

Themis Palpanas, Michail Vlachos, Eamonn J. Keogh, and Dimitrios Gunopulos. 2008. Streaming Time Series Summarization Using User-Defined Amnesic Functions. *IEEE Transactions on Knowledge and Data Engineering* 20 (2008), 992–1006. Issue 7.

Jia-Yu Pan, Hyung-Jeong Yang, and Christos Faloutsos. 2004. MMSS: Multi-Modal Story-Oriented Video Summarization. In *Proceedings of the 4th IEEE International Conference on Data Mining (ICDM), Brighton, UK*. 491–494.

Spiros Papadimitriou, Jimeng Sun, Christos Faloutsos, and Philip S. Yu. 2008. Hierarchical, Parameter-Free Community Discovery. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), Antwerp, Belgium*.

Jian Pei, Daxin Jiang, and Aidong Zhang. 2005. On Mining Cross-Graph Quasi-Cliques. In *Proceedings of the 11th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Chicago, IL*. 228–238.

Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. DeepWalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), New York, NY*. ACM, 701–710.

B. Aditya Prakash, Jilles Vreeken, and Christos Faloutsos. 2012. Spotting Culprits in Epidemics: How many and Which ones?. In *Proceedings of the 12th IEEE International Conference on Data Mining (ICDM), Brussels, Belgium*. IEEE.

Manish Purohit, B Aditya Prakash, Chanhyun Kang, Yao Zhang, and VS Subrahmanian. 2014. Fast Influence-Based Coarsening For Large Networks. In *Proceedings of the 20th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), New York, NY*. ACM, 1296–1305.

Tai Qin and Karl Rohe. 2013. Regularized Spectral Clustering Under the Degree-Corrected Stochastic Blockmodel. In *Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS), South Lake Tahoe, UT*. MIT Press, 3120–3128.

Qiang Qu, Siyuan Liu, Christian S. Jensen, Feida Zhu, and Christos Faloutsos. 2014. Interestingness-Driven Diffusion Process Summarization in Dynamic Networks. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), Nancy, France*. 597–613.

Davood Rafiei and Stephen Curial. 2005. Effectively Visualizing Large Networks Through Sampling. In *16th IEEE Visualization Conference (VIS 2005), Minneapolis, MN, USA*. 48.

Sriram Raghavan and Hector Garcia-Molina. 2003. Representing Web Graphs. In *Proceedings of the 19th International Conference on Data Engineering (ICDE), Bangalore, India*. IEEE, 405–416.

Matteo Riondato, David García-Soriano, and Francesco Bonchi. 2014. Graph Summarization With Quality Guarantees. In *Proceedings of the 14th IEEE International Conference on Data Mining (ICDM), Shenzhen, China*. IEEE, 947–952.

Karl Rohe, Sourav Chatterjee, and Bin Yu. 2011. Spectral Clustering and the High-Dimensional Stochastic Blockmodel. *The Annals of Statistics* 39, 4 (2011), 1878–1915.

Ryan Rossi, Brian Gallagher, Jennifer Neville, and Keith Henderson. 2012. Role-dynamics: fast mining of large dynamic networks. In *Proceedings of the 21st International Conference on World Wide Web (WWW), Lyon, France (WWW '12 Companion)*. ACM, 997–1006.

Martin Rosvall and Carl T. Bergstrom. 2007. An Information-Theoretic Framework for Resolving Community Structure in Complex Networks. *Proceedings of the National Academy of Sciences* 104, 18 (2007), 7327–7331.

Randolf Rotta and Andreas Noack. 2011. Multilevel Local Search Algorithms for Modularity Clustering. *Journal of Experimental Algorithmics (JEA)* 16 (2011).

Barna Saha and Pabitra Mitra. 2007. Dynamic Algorithm for Graph Clustering Using Minimum Cut Tree.. In *SDM*. SIAM, 581–586.

Neil Shah, Alex Beutel, Brian Gallagher, and Christos Faloutsos. 2014. Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective. In *Proceedings of the 14th IEEE International Conference on Data Mining (ICDM), Shenzhen, China*. IEEE.

Neil Shah, Danai Koutra, Tianmin Zou, Brian Gallagher, and Christos Faloutsos. 2015. TimeCrunch: Interpretable Dynamic Graph Summarization. In *Proceedings of the 21st ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Sydney, Australia*.

Umang Sharan and Jennifer Neville. 2008. Temporal-Relational Classifiers for Prediction in Evolving Domains. In *Proceedings of the 8th IEEE International Conference on Data Mining (ICDM), Pisa, Italy*. 540–549.

Zeqian Shen, K-L Ma, and Tina Eliassi-Rad. 2006. Visual Analysis of Large Heterogeneous Social Networks by Semantic and Structural Abstraction. *IEEE Transactions on Visualization and Computer Graphics* 12, 6 (2006), 1427–1439.

Lei Shi, Hanghang Tong, Jie Tang, and Chuang Lin. VEGAS: Visual InfluEnce GrAph Summarization on Citation Networks. *IEEE Transactions on Knowledge and Data Engineering* 27, 12 (????), 3417–3431.

Ben Shneiderman. 2008. Extreme Visualization: Squeezing a Billion Records into a Million Pixels. In *Proceedings of the 2008 ACM International Conference on Management of Data (SIGMOD), Vancouver, BC*.

Mahsa Shoaran, Alex Thomo, and Jens H Weber-Jahnke. 2013. Zero-Knowledge Private Graph Summarization. In *IEEE International Conference on Big Data*. IEEE, 597–605.

Koen Smets and Jilles Vreeken. 2011. The Odd One Out: Identifying and Characterising Anomalies. In *Proceedings of the 11th SIAM International Conference on Data Mining (SDM), Mesa, AZ*. SIAM, 804–815.

Sucheta Soundarajan, Acar Tamersoy, Elias B. Khalil, Tina Eliassi-Rad, Duen Horng Chau, Brian Gallagher, and Kevin Roundy. 2016. Generating Graph Snapshots from Streaming Edge Data. In *Proceedings of the 25th International Conference Companion on World Wide Web*. 109–110.

Daniel A. Spielman and Nikhil Srivastava. 2008. Graph Sparsification by Effective Resistances. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. ACM, 563–568.

Jimeng Sun, Christos Faloutsos, Spiros Papadimitriou, and Philip S Yu. 2007. GraphScope: Parameter-Free Mining of Large Time-Evolving Graphs. In *Proceedings of the 13th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Jose, CA*. ACM, 687–696.

Jimeng Sun, Charalampos E. Tsourakakis, Evan Hoke, Christos Faloutsos, and Tina Eliassi-Rad. 2008. Two heads better than one: pattern discovery in time-evolving multi-aspect data. *Data Min. Knowl. Discov.* 17, 1 (2008), 111–128.

Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. LINE: Large-scale Information Network Embedding. In *Proceedings of the 24th International Conference on World Wide Web (WWW), Florence, Italy*. 1067–1077.

Nikolaj Tatti and Jilles Vreeken. 2012. The Long and the Short of It: Summarizing Event Sequences with Serial Episodes. In *Proceedings of the 18th ACM International*

Conference on Knowledge Discovery and Data Mining (SIGKDD), Beijing, China. ACM.

Lini T. Thomas, Satyanarayana R. Valluri, and Kamalakar Karlapalem. 2010. MARGIN: Maximal Frequent Subgraph Mining. *ACM Transactions on Knowledge Discovery From Data* 4, 3 (2010), 10:1–10:42.

Yuanyuan Tian, Richard A Hankins, and Jignesh M Patel. 2008. Efficient Aggregation for Graph Summarization. In *Proceedings of the 2008 ACM International Conference on Management of Data (SIGMOD), Vancouver, BC*. ACM, 567–580.

Hannu Toivonen, Fang Zhou, Aleksi Hartikainen, and Atte Hinkka. 2011. Compression of Weighted Graphs. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Diego, CA*. 965–973.

Hanghang Tong, Christos Faloutsos, and Jia-Yu Pan. 2006. Fast Random Walk with Restart and Its Applications. In *ICDM*. 613–622.

Jilles Vreeken, Matthijs van Leeuwen, and Arno Siebes. 2011. KRIMP: Mining Itemsets that Compress. *Data Mining and Knowledge Discovery* 23, 1 (2011), 169–214.

Bianca Wackersreuther, Peter Wackersreuther, Annahita Oswald, Christian Böhm, and Karsten M. Borgwardt. 2010. Frequent Subgraph Discovery in Dynamic Networks. In *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*. ACM, 155–162.

Chao Wang and Srinivasan Parthasarathy. 2006. Summarizing itemset patterns using probabilistic models. In *Proceedings of the 12th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Philadelphia, PA*. 730–735.

Daixin Wang, Peng Cui, and Wenwu Zhu. 2016. Structural Deep Network Embedding. In *Proceedings of the 22nd ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), San Francisco, CA*. KDD.

Jianyong Wang and George Karypis. 2004. SUMMARY: Efficiently Summarizing Transactions for Clustering. In *Proceedings of the 4th IEEE International Conference on Data Mining (ICDM), Brighton, UK*. 241–248.

Ye Wu, Zhinong Zhong, Wei Xiong, and Ning Jing. 2014. Graph Summarization for Attributed Graphs. In *International Conference on Information Science, Electronics and Electrical Engineering (ISEEE)*. IEEE, 503–507.

Y. Xiang, R. Jin, D. Fuhry, and F. Dragan. 2010. Summarizing transactional databases with overlapped hyperrectangles. *Data Mining and Knowledge Discovery* (2010).

Kevin S Xu, Mark Kliger, and Alfred O Hero III. 2011. Tracking Communities in Dynamic Social Networks. In *Proceedings of the 4th International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction (SBP'11)*. Springer, 219–226.

Zhiqiang Xu, Yiping Ke, Yi Wang, Hong Cheng, and James Cheng. 2012. A Model-based Approach to Attributed Graph Clustering. In *Proceedings of the 2012 ACM International Conference on Management of Data (SIGMOD), Scottsdale, AZ*. ACM.

Xifeng Yan, Hong Cheng, Jiawei Han, and Dong Xin. 2005. Summarizing itemset patterns: a profile-based approach. In *Proceedings of the 11th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD), Chicago, IL*. 314–323.

Jaewon Yang and Jure Leskovec. 2013. Overlapping Community Detection at Scale: A Nonnegative Matrix Factorization Approach. In *Proceeding of the 6th ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, 587–596.

Jaewon Yang, Julian McAuley, and Jure Leskovec. 2013. Community Detection in Networks with Node Attributes. In *Proceedings of the 14th IEEE International Conference on Data Mining (ICDM), Dallas, Texas*. IEEE, 1151–1156.

Jinguo You, Qiuping Pan, Wei Shi, Zhipeng Zhang, and Jianhua Hu. 2013. Towards Graph Summary and Aggregation: A Survey. In *Social Media Retrieval and Mining*.

Springer, 3–12.
Ning Zhang, Yuanyuan Tian, and Jignesh M. Patel. 2010. Discovery-driven Graph Summarization. In *Proceedings of the 26th International Conference on Data Engineering (ICDE), Long Beach, CA*. 880–891.
Yang Zhou, Hong Cheng, and Jeffrey Xu Yu. 2009. Graph clustering based on structural/attribute similarities. *Proceedings of the VLDB Endowment* 2, 1 (2009), 718–729.
Linhong Zhu, Majid Ghasemi-Gol, Pedro Szekely, Aram Galstyan, and Craig A Knoblock. 2016. Unsupervised Entity Resolution on Multi-type Graphs. In *International Semantic Web Conference*. Springer, 649–667.